

2016

Point process modeling and estimation: advances in the analysis of dynamic neural spiking data

<https://hdl.handle.net/2144/17719>

Downloaded from DSpace Repository, DSpace Institution's institutional repository

BOSTON UNIVERSITY
GRADUATE SCHOOL OF ARTS AND SCIENCES

Dissertation

**POINT PROCESS MODELING AND ESTIMATION: ADVANCES IN
THE ANALYSIS OF DYNAMIC NEURAL SPIKING DATA**

by

XINYI DENG

B.S., American University, 2011

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

2016

© 2016 by
XINYI DENG
All rights reserved

Approved by

First Reader

Uri T. Eden, PhD
Associate Professor of Mathematics and Statistics

Second Reader

Eric D. Kolaczyk
Professor of Mathematics and Statistics

Third Reader

Loren M. Frank
Professor of Physiology, UCSF

“That’s the effect of living backwards,” the Queen said kindly: “it always makes one a little giddy at first—”

“Living backwards!” Alice repeated in great astonishment. “I never heard of such a thing!”

“—but there’s one great advantage in it, that one’s memory works both ways.”

“I’m sure *mine* only works one way,” Alice remarked. “I can’t remember things before they happen.”

“It’s a poor sort of memory that only works backwards,” the Queen remarked.

—Lewis Carroll, *Through the Looking-Glass*

Acknowledgments

Even in my wildest dream, I could not have foreseen the profound impact the decision to join the graduate program in Statistics at Boston University four and half years ago would have on my life. It opened a door to a world at the intersection of statistics and neuroscience that I didn't know existed. I owe my gratitude to all those people who have made my graduate experience one that I will always cherish.

My deepest gratitude is to my thesis advisor, Prof. Uri Eden, who in the past four years has been nothing but supportive and kind. I have been incredibly lucky to work with such an amazing mentor who takes all aspects of his job extremely seriously, an engaging teacher with penetrating clarity, and a brilliant statistician with immense dedication to his craft. My apprenticeship with Prof. Eden has been truly inspirational. I hope that one day I would become as good an advisor to my students as Prof. Eden has been to me.

I am also deeply grateful to Prof. Loren Frank. My interactions with Prof. Frank and his group during my summer visits to San Francisco and our numerous Skype meetings have provided me with not only the knowledge necessary to undertake interdisciplinary research but also the knowledge that doing good science is such a fulfilling and joyous thing. Our collaboration is one of the most fun and exciting parts of my graduate experience, due in no small part to Prof. Frank's intellectual brilliance and great sense of optimism.

I would also like to extend my gratitude to the members of my thesis committee. Prof. Eric Kolaczyk has been instrumental in helping me chart an academic path and guiding me through different stages of my graduate career. Prof. Mark Kramer has contributed a great deal to create an open and friendly environment within the group and I have many fond memories of brainstorming research ideas at our Friday group meetings. Prof. Luis Carvalho is one of the best teachers that I have had whose

lectures have often benefited my research.

I would also like to acknowledge Prof. Murad Taqqu and Prof. Mamikon Ginovyan whose invaluable lectures on Probability Theory and Estimation Theory transformed my statistical thinking and continue to provide inspiration.

I am also indebted to my many colleagues and friends. Particularly, I must thank Daniel Liu, who as a close collaborator and a good friend, was always willing to help me with experimental data, to provide me with any neuroscience knowledge that I lack, and to encourage me with the occasional pep talks. Thanks to Jessica Curtis and Dara Gold for making our little office a sanctuary and helping me cope with stress. I am also grateful to Mikio Aoi, Wei Tang, and Louis-Emmanuel Martinet for many intellectually stimulating conversations over Thai food on “Noodle Friday”. Thanks to Lijun Peng, Yaonan Zhang, Han Wang, Xinyu Kang, Shuyang Bai, Xinying Mu, and Yujia Zhou for their help and support. Thanks also to Zarek Brot-Goldberg and Huadi Zhang for always being there for me; I greatly value our friendship.

During this period, I have had the privilege to collaborate with some extraordinary researchers including Kenneth Kay, Dr. Kensuke Arai, Dr. Mattias Karlsson, Dr. Rose Faghieh, Dr. Riccardo Barbieri, Dr. Angelique Paulk, Dr. Wael Asaad, Dr. Emery Brown, Dr. Darin Dougherty, Dr. Alik Widge, and Dr. Emad Eskandar.

I am grateful to Prof. Tasso Kaper and all the staff in the Department of Mathematics and Statistics for their various forms of support during my graduate study.

Finally, I offer my endless appreciation to my parents—my father for his quiet calm and my mother for having taught me the importance of independent thoughts, for them helping me through trying times, and mostly, for their unconditional love and support.

**POINT PROCESS MODELING AND ESTIMATION: ADVANCES IN
THE ANALYSIS OF DYNAMIC NEURAL SPIKING DATA**

XINYI DENG

Boston University, Graduate School of Arts and Sciences, 2016

Major Professor: Uri T. Eden, PhD

Associate Professor of Mathematics and Statistics

ABSTRACT

A common interest of scientists in many fields is to understand the relationship between the dynamics of a physical system and the occurrences of discrete events within such physical system. Seismologists study the connection between mechanical vibrations of the Earth and the occurrences of earthquakes so that future earthquakes can be better predicted. Astrophysicists study the association between the oscillating energy of celestial regions and the emission of photons to learn the Universe's various objects and their interactions. Neuroscientists study the link between behavior and the millisecond-timescale spike patterns of neurons to understand higher brain functions.

Such relationships can often be formulated within the framework of state-space models with point process observations. The basic idea is that the dynamics of the physical systems are driven by the dynamics of some stochastic state variables and the discrete events we observe in an interval are noisy observations with distributions determined by the state variables. This thesis proposes several new methodological developments that advance the framework of state-space models with point process observations at the intersection of statistics and neuroscience. In particular, we develop new methods 1) to characterize the rhythmic spiking activity using history-dependent structure, 2) to model population spike activity using marked point process models, 3) to allow for real-time decision making, and 4) to take into account the need

for dimensionality reduction for high-dimensional state and observation processes.

We applied these methods to a novel problem of tracking rhythmic dynamics in the spiking of neurons in the subthalamic nucleus of Parkinson's patients with the goal of optimizing placement of deep brain stimulation electrodes. We developed a decoding algorithm that can make decision in real-time (for example, to stimulate the neurons or not) based on various sources of information present in population spiking data. Lastly, we proposed a general three-step paradigm that allows us to relate behavioral outcomes of various tasks to simultaneously recorded neural activity across multiple brain areas, which is a step towards closed-loop therapies for psychological diseases using real-time neural stimulation. These methods are suitable for real-time implementation for content-based feedback experiments.

Contents

1	Introduction	1
1.1	Point Process Models	1
1.1.1	Theoretical basics	2
1.1.2	Point process models in seismology: earthquakes	4
1.1.3	Point process models in astronomy: gamma-ray bursts	8
1.1.4	Point process models in neuroscience: spikes	11
1.2	State-space Models	13
1.2.1	Theoretical basics	13
1.2.2	State-space models in the analysis of dynamic neural spiking data	14
1.3	Goals of the thesis	15
2	A Point Process Approach to Identifying and Tracking Transitions in Neural Spiking Dynamics in the Subthalamic Nucleus of Parkinson’s Patients	19
2.1	Introduction	19
2.2	Methods	23
2.2.1	Detecting changes in history-dependent firing patterns	24
2.2.2	State-space smoothing algorithm to track the change through time	29
2.3	Results	31
2.3.1	Preliminary analysis	31
2.3.2	Identifying significant changes in spiking dynamics— simulation study	33

2.3.3	Identifying significant changes in spiking dynamics—STN data	35
2.3.4	Tracking changes in spiking dynamics—simulation study . . .	38
2.3.5	Tracking changes in spiking dynamics—STN data	41
2.4	Discussion	43
3	Clusterless Decoding of Position From Multiunit Activity Using A Marked Point Process Filter	53
3.1	Introduction	53
3.2	Methods	56
3.3	Simulation Study	60
3.3.1	Data simulation	61
3.3.2	Decoding results	63
3.3.3	Goodness-of-fit analysis	66
3.4	An Application to Position Decoding from Multiunit Activity in Rat Hippocampus	69
3.5	Discussion	73
3.6	Appendix	76
3.6.1	Hippocampal data collection and preprocessing	76
3.6.2	Gaussian approximation to the posterior density	77
3.6.3	Sequential Monte Carlo decoding algorithm	78
4	Rapid Classification of Hippocampal Replay Content for Real-time Applications	80
4.1	Introduction	80
4.2	Algorithm Development	82
4.2.1	Discrete state point process filters	83
4.2.2	Algorithm discussion	87
4.3	Data Analyses	89

4.3.1	Description of the experimental data	90
4.3.2	Analysis of individual replay events: sources of information for decoding replay position	91
4.3.3	Analysis of individual replay events: decoding replay position	93
4.3.4	Analysis of individual replay events: sources of information for decoding a decision state	94
4.3.5	Analysis of individual replay events: decoding a decision state	98
4.3.6	Analysis of an entire recording epoch: decoding a decision state	101
4.4	Discussion and Conclusion	102
4.5	Appendix	105
4.5.1	Bootstrap confidence bounds for the estimated joint-mark in- tensity	105
4.5.2	Incorporating model uncertainty into decoding algorithm	106
5	Estimating a Dynamic State to Relate Neural Spiking Activity to Behavioral Signals during Cognitive Tasks	114
5.1	Introduction	114
5.2	Cognitive State-space Decoding Paradigm	116
5.2.1	Model framework	116
5.2.2	Estimating state dynamics from behavior	118
5.2.3	Neural encoding	118
5.2.4	Neural decoding	119
5.3	Application: Decoding Learning State from Spiking Activity in Mon- key PFC and Caudate	120
5.3.1	Experimental data	120
5.3.2	Estimate state dynamics from behavior	120
5.3.3	Neural encoding	122

5.3.4	Neural decoding	126
5.4	Discussion	127
6	Conclusions	130
6.1	Contributions of the Thesis	130
6.2	Future Research	132
	References	134
	Curriculum Vitae	148

List of Figures

1.1	Photon counts of a particular GRB event, <i>Trigger 0551</i> , as a function of time. (Reproduced from Dr. Jessie Henshaw’s personal archive: http://www.synapse9.com/ .)	9
1.2	Schematic outline of the thesis.	17
2.1	(A) Raster and peri-stimulus histogram (PSTH) of 3.2 second recordings from a single neuron before and after onset of hand movement (time 0) for all trials. (B) Trial-averaged autocorrelation functions. The red lines denote pointwise 95% significance bounds. (C) Trial-averaged spectrograms of spike trains as a function of time relative to movement onset on the x-axis and frequency between 0–100 Hz on the y-axis.	31

2.2 (A) Estimates and 95% confidence bounds for estimated transition times for the start and end points of the movement spiking state for simulated spike trains of four neurons under the two-state model. Actual transitions occurred at -400 ms and 400 ms for start and end transitions, respectively. (B) 95% confidence bounds for the estimated history dependent modulation parameters as a function of lag and the true parameter values for simulated spike data from a single neuron. Cyan dashed line and magenta dotted line represent true modulation during non-movement state and during movement state, respectively. The transparent regions in cyan and magenta correspond to the estimated 95% confidence regions. These regions are narrow and tend to contain the true value at every lag. (C) K-S plot comparing empirical and model CDFs of rescaled simulated ISIs from estimated point process model. (D) Power of the maximum likelihood ratio test as a function of the modulation of the history parameters at 30 ms and 50 ms lags between the non-movement and movement states. The test detects as little as 2.4% modulation in these parameters with 82.5% probability. 48

2·3	<p>(A) Estimates and 95% confidence bounds for transition points between non-movement and movement states for four representative neurons.</p> <p>(B) Estimated history dependent modulation parameters as a function of lag with 95% confidence bounds for a single representative neuron. Cyan dashed line and magenta dotted line represent modulation estimates during non-movement state and during movement state, respectively. Transparent regions represent corresponding the 95% confidence regions.</p> <p>(C) K-S plot comparing empirical and model CDFs of rescaled ISIs from estimated point process model.</p> <p>(D) P-values for the maximum likelihood ratio test for significant changes in modulation between states for all 22 cells, in descending order.</p>	49
2·4	<p>Histograms of ratio of estimated model parameters for baseline intensity and modulation at lags of 2, 30, and 50 ms across all 22 neurons analyzed. A value of 1 indicates no change in modulation between the two states, a value above 1 indicates an increased estimate in the movement state, relative to the non-movement state, and a value below 1 indicates a decreased estimate in the movement state.</p>	50

2·5	<p>(A) Estimated trajectory and uncertainty of baseline firing rate parameter for simulated spiking data. Blue line represents the estimates and red lines denote 95% confidence regions at each time. Black line represents the true constant baseline firing rate parameter used to simulate the data. (B) Estimated and true values of history dependent modulation of firing intensity as a function of time relative to movement onset on the x-axis and lag between 10–100 ms on the y-axis. (C) Standard deviations for parameter estimates for baseline firing rate, and modulation at lag 2 ms , 30 ms and 50 ms lags. In each case, estimated standard deviations are small compared to estimated mean parameter values, indicating high confidence in estimates. (D) K-S plot comparing empirical and model CDFs of rescaled simulated ISIs from estimated point process model.</p>	51
2·6	<p>(A) Estimated trajectory and uncertainty of baseline firing rate parameter for observed spiking data from a representative STN neuron. Blue line represents the estimates and red lines denote 95% confidence regions at each time. (B) Estimated values of history dependent modulation of firing intensity as a function of time relative to movement onset on the x-axis and lag between 10–100 ms on the y-axis. Blue regions at beginning and end of trial from 20 to 35 ms lags indicate regions of negative modulation and red regions at beginning and end of trial from 40 to 60 ms lags indicate regions positive modulation. (C) Standard deviations for parameter estimates for baseline firing rate, and modulation at lag 2 ms , 30 ms and 50 ms. (D) K-S plot comparing empirical and model CDFs of rescaled ISIs from estimated point process model.</p>	52

3.1 (A) Simulated trajectory of the animal running back and forth on a linear track. (B) One second unsorted spike train with marks from two simulated neurons. (C) One second raster plot of spikes without marks where the spikes were sorted via a linear discriminant function on their mark values. (D) True joint mark intensity of the two simulated neurons as a function of linear position on the y-axis and mark value on the x-axis. The place fields center at -1.5 and 1.5 respectively, with a 0.1 variance. The mark spaces are one-dimensional with mean of 10 and 13 respectively, and standard deviation of 2. (E) True rate of each of the two simulated neurons as a function of linear position. It represents the place field of each neuron. 62

3.2 Decoding results for two simulated neurons. (A) Posterior density of the animal’s linear position using the clusterless decoding algorithm. The blue line represents the actual position of the animal. The red represents the posterior density at each time step. (B) Posterior density of the animal’s linear position using the decoding with spike-sorting. (C) Clusterless decoding results zoomed in at time between 400 and 600 ms, showing the multimodality of the posterior around 430 ms. (D) Decoding with spike-sorting results zoomed in at time between 400 and 600 ms. 63

3.3	<p>Quality of fit comparison between the two decoding algorithms. (A) Root-mean-squared error (rMSE) between the true positions of the animal and their estimated values averaged across 100 trials as a function of the overlap between mark spaces of two simulated neurons. Error bars represent 2 standard deviations from the mean rMSE. The dotted line represents the performance of clusterless decoding method. The solid line represents the performance of decoding with spike-sorting method. (B) Fraction of time that the true position values were covered by the 99% highest posterior density (HPD) region averaged across 100 trials as a function of the overlap between mark spaces of two simulated neurons. Error bars represent 2 standard deviations from the mean coverage probability. The dotted line represents the performance of clusterless decoding method. The solid line represents the performance of decoding with spike-sorting method.</p>	67
3.4	<p>Decoding results for hippocampus data. (A) Posterior density of the rat's linear position using the clusterless decoding algorithm. The blue line represents the actual position of the rat. The red represents the posterior density at each time step. The posterior density computed through the clusterless decoding algorithm has an rMSE of 14.3 cm and remains within the 99% HPD region 74.25% of the time. (B) Posterior density of the rat's linear position using decoding with manual spike-sorting. The posterior density computed by first spike-sorting and then decoding has a rMSE of 26.0 cm and remains within the 99% HPD region 70.76% of the time.</p>	72

4.1	<p>Linearization and encoding models for position. (A) 2D representation of an M-maze. The black triangle represents an animal whose head is oriented towards the choice-point. (B) Initial condition of the state variable, linearized position $x(t)$. Here we assume a replay event is equally likely to start at any position on the maze. (C) State transition probability matrix $p(x_k x_{k-1})$ of transition probabilities from every possible replay position at the prior time step to each value at current time. Here we posit that the dynamics of the replayed trajectory are similar to those of the actual movement and use the empirical movement data to compute the state transition model. (D) Observation model for an example tetraode. Here we plot the pairwise marginalization of the 5-dimension position-mark model. Panel at the intersection of each row and column plots the estimated joint mark intensity function marginalized over the row and column linearized position or tetraode channel.</p>	107
4.2	<p>Decoding replay trajectory from multiunit spiking activity using a marked point process filter. (A, B, C, D): <i>Left</i> Unsorted ensemble spiking activity on each tetraode as a function of replay time. <i>Right</i> Replay trajectories of four examples decoded using clusterless methods where the heat plot shows the estimated posterior density at each time step. <i>Top</i> Schematic representation of the decoding results where the yellow arrowed line illustrates the evolution of the replay trajectory.</p>	108

4.3	<p>Examples of the three sources of information contributing to the decision state decoder. As a visualization aid, we plot sources of information that are consistent with those during actual outbound movement in shades of blue and sources of information that are consistent with those during actual inbound movement in shades of red. (A, B, C, D): <i>Left</i> Initial condition of replay position conditioned on the decision state. <i>Center</i> State transition model of replay position conditioned on the decision state. A slice (when the animal is at a position on the left arm shown as the black dot) of the state transition matrix is plotted here <i>Right</i> Estimated observation or likelihood model of joint mark intensity conditioned on the decision state.</p>	109
4.4	<p>Decoding a discrete, fixed decision state. (A, B, C, D): Probability of a decision state of four example replay events as a function of time. Probability of the replay event representing an “outbound and forward” path, an “outbound and reverse” path, an “inbound and forward” path, or an “inbound and reverse” path is plotted in darker blue, lighter blue, darker red, and lighter red, respectively <i>Top</i> Schematic representation of the decoding results where the yellow arrowed line represents the temporal evolution of the replay trajectory and the black triangle represents the replayed orientation of the rat at the beginning of the event.</p>	110
4.5	<p>Interpreting the decision state filter as a set of weighted, parallel marked point process filters. (A, B, C, D) Decoded replay trajectory conditioning on each of the decision states for four example replay events where the heat plot shows the estimated posterior density at each time step.</p>	111

4.6	Relative frequency histogram of decision state as a function of time into a replay event, with four different thresholds on $\Pr(I)$. Events that can be classified are plotted in green. Events whose decision state probabilities did not pass the threshold and cannot be classified are plotted in black.	112
4.7	Decision state decoding results for two recording epochs on the same experiment day. Summary of decision state, by lap, with a threshold of 0.9. Replay event classified as an “outbound, forward” path, an “outbound, reverse” path, an “inbound, forward” path, or an “inbound, reverse” path is plotted in darker blue, lighter blue, darker red, and lighter red, respectively. On the axis, left and outbound lap, left and inbound lap, right and outbound lap, and right and inbound lap are denoted on the x-axis as “oL”, “iL”, “oR”, and “iR”. (A) A novel track. (B) A familiar track.	113
5.1	Schematic representation of the general three-model cognitive state-space paradigm (<i>Lower</i>) in comparison with the classical two-model state-space paradigm (<i>Upper</i>).	117
5.2	Two examples of the EM algorithm applied in the analysis of the state dynamics in a rule-switching task. The correct and incorrect responses are shown, respectively, by black and gray marks above the panels. The probability of a correct response occurring by chance is 0.25 in horizontal line. Solid blue lines are the learning curve estimates, and the dotted blue lines are the associated 95% confidence intervals. . . .	121

- 5.3 Four examples of the estimated conditional intensity with 95% confidence bounds as a function of the temporal features in the task. In each panel, blue and red solid line is the estimated conditional intensity for the “not learned” and “learned” trials, respectively. The vertical lines in cyan, green, yellow and black are the time for picture cue, go cue, feedback, and start of the inter-trial interval, respectively. . . . 124
- 5.4 *Left:* An example Kolmogorov-Smirnov (KS) plot of time-rescaled inter-spike intervals for the estimated point process neural encoding model. It plots the empirical cumulative distribution of the rescaled ISIs against the theoretical cumulative distribution of the Exponential(1) distribution. *Right:* An example receiver-operating characteristic (ROC) curve. It plots sensitivity of the cut-off, the probability of rejecting the null hypothesis when it is false, versus significance level, the probability of rejecting a null hypothesis when it is true. ROC curves using decoding results from the neural ensemble in PFC and caudate are plotted in solid blue lines and dashed red line, respectively. 125

List of Abbreviations

ACF	Autocorrelation Function
CDF	Cumulative Distribution Function
EEG	Electroencephalography
EM	Expectation-Maximization
GLM	Generalized Linear Model
HPD	Highest Posterior Density
ISI	Inter-spike Interval
K-S plot	Kolmogorov-Smirnov plot
LFP	Local Field Potentials
MEG	Magnetoencephalography
PFC	Prefrontal Cortex
PSTH	Peri-stimulus Histogram
rMSE	root-Mean-Squared Error
ROC curve	Receiver-operating Characteristic curve
SSPPF	Stochastic State Point Process Filter
STN	Subthalamic Nucleus
SWR	Sharp-wave Ripple

Chapter 1

Introduction

Point processes, stochastic processes whose realizations consist of point events in time or space, arise naturally in many fields of application, and because of their simple qualitative structure and tractable but broad special cases, have been extensively studied (Cox and Isham, 1980; Karr, 1991).

Descriptions of point process data have existed since ancient time. Our ancestors recorded speculation about the distribution of stars in a sky once thought spherical and occurrences of floods, earthquakes and other natural events that exhibit no evident periodic structure (Karr, 1991). By contrast, the modern mathematical theory for point processes is a much more recent development (Palm, 1943; Jensen, 1948; Wold, 1948; Cox, 1955; Khintchine, 1956a; Khintchine, 1956b).

Since, there has developed on the one hand a general, powerful and elegant collection of mathematical theory on point processes, and on the other hand diverse applications to fields such as biology, geography, meteorology, physics, operation research and engineering. It confirms that important physical problems lead to the most challenging and interesting mathematical questions (Karr, 1991).

1.1 Point Process Models

In this section, we first introduce some of the most important ideas needed to develop probabilistic models in the point processes context. We then describe recent work on some statistical methods with applications, particularly in seismology, astrophysics,

and neuroscience. One of the principal goals of this section is to bring out the universality of statistical techniques by examples from these three fields—we show the same methods playing central roles in the analysis of point process data from quite disparate fields.

1.1.1 Theoretical basics

In this subsection, we introduce some of the most important ideas needed to develop probabilistic models in the point processes context. Many texts on point processes can be referred to for further details: Bremaud (1981) and Daley and Vere-Jones (2003) establish the background theory, Snyder (1975), Cox and Isham (1980) and Snyder and Miller (1995) bring modelling aspects and some statistical applications, and Karr (1991) (my personal favorite) offers an enjoyable read on the extensiveness and flexibility of the point process framework, with more emphasis on probability than statistics.

Just like all roads eventually lead to Rome, mathematical thinking from different starting points (for example, counting theory, renewal theory, and Martingale theory) can lead to the point process framework. Here we follow the development laid out by Daley and Vere-Jones (Daley and Vere-Jones, 2003; Vere-Jones, 1995).

The point process framework can be used to describe observations whose occurrences can be treated as time instants, $N(t)$, with which other variables can be associated. If we choose to focus on the sequence of time intervals between events, or increments, then we need to first take into account the time which has already elapsed since the occurrence of the last event.

A special case is when the intervals are presumed independent, forming a renewal process. Then the conditional distribution of the remaining lifetime given the time

already elapsed can be represented in the form of the hazard function $h(t)$,

$$h(t)dt = \Pr(\text{life} < t + dt | \text{life} > t), \quad (1.1)$$

where lifetimes are the interval lengths.

If the life distribution has density $f(t)$ and distribution function $F(t)$, then

$$h(t) = \frac{f(t)}{1 - F(t)}. \quad (1.2)$$

Then,

$$P(\text{life} > t + \tau | \text{life} > t) = \exp\left\{-\int_t^{t+\tau} h(u)du\right\}. \quad (1.3)$$

In the general case, the hazard function $h(t)$ has to be conditioned, not only on the time since the last event, but on any additional information concerning the past history H_t . The hazard function $h(t)$ then becomes the conditional probability intensity, or conditional intensity function, $\lambda(t|H_t)$:

$$\lambda(t|H_t)dt = E[dN(t)|\text{past history up to time } t]. \quad (1.4)$$

It is possible to define the past history more generally to include: information about additional variables associated with each event, or marks, each of which can be a vector; information about external variables; processes evolving in time in parallel with the interested point process; etc.

The conditional intensity function plays a pivotal role in developing modeling and estimation methods for point processes. Some of its key properties are:

1. A knowledge of the conditional intensity function completely defines the probability structure of the process (Liptser and Shiryaev, 1978).
2. The log-likelihood of a realization (t_1, \dots, t_N) in the time interval $[0, T]$ is given

by

$$\log L(t_1, \dots, t_N) = \sum_i^N \log \lambda(t_i) - \int_0^T \lambda(u) du. \quad (1.5)$$

3. If T^* denotes the time to the next event,

$$P(T^* > \tau) = \exp\left\{-\int_0^\tau \lambda^*(u) du\right\}, \quad (1.6)$$

where $\lambda^*(u)$ is obtained by extrapolating $\lambda(t|H_t)$ forward from t to $t+u$ assuming no event occurs in $(t, t+u)$.

4. The differences $dN(t) - \lambda(t|H_t)dt$ integrate to form a continuous time Martingale.

Important consequences follow:

1. Under some regularity conditions, the full machinery of likelihood methods can be used with the conditional intensity models, including estimation methods, hypothesis testing methods, likelihood-based model selection criteria, etc.
2. The conditional intensity models can be used as a basis for simulations, in particular, forward simulations given a past history.
3. The conditional intensity models provides a link to the Martingale properties.

Determining how the conditional intensity depends on past history and external variables require not only mathematical knowledge, but also insights into the physical systems in which we are interested. To emphasize this point, we discuss some applications of the point process framework in the following subsections.

1.1.2 Point process models in seismology: earthquakes

Seismology is a science based on data called “seismograms”. A seismogram is a time series of the displacement, velocity, or acceleration experienced by a particle of

the Earth and is a record of mechanical vibrations of the Earth (Aki and Richards, 1980). By carefully arranging seismometers of various dimensions in spatial arrays, an earthquake signal can be seen as travelling and changing shape (Brillinger, 1988b). Additional information about individual earthquakes, such as occurrence times, locations, sizes and other characteristics, is documented in “catalogs”.

Seismologists are interested in structural questions such as: How should seismometers be laid out in a network (Udias, 1989)? Is activity on different faults associated (Sanders, 1993; Hill et al., 1993; Wesson et al., 2003)? Are layer boundaries flat or bumpy (Pulliam and Stark, 1993)? Do smaller or larger earthquakes provide a more accurate and detailed description of faults (Schaff et al., 2002; Rubinstein and Beroza, 2007)? However, the problems of forecasting earthquakes and earthquake risk have always remained the focus of attention, especially with the growth of major cities in hazard-prone areas and the public anxiety associated with risks to critical facilities such as nuclear reactors (Vere-Jones, 1995). In seismology, the phrase *earthquake forecasting* means, specifically, the assignment of probabilities of earthquakes, over temporal intervals ranging from decades to centuries, occurring in broader space-time-magnitude regions (Hough, 2010). It is contrasted with the related phrase *earthquake prediction*, the identification of a meaningfully small geographic region and time window in which a major earthquake will occur with very high probability.

A broad range of point process models have been employed by statisticians researching seismic activity with applications to earthquake forecasting: The sequence of times $\{t_i\}$ of earthquake occurrences in a given region can be viewed as a realization of a point process. It becomes a marked point process $\{t_i, M_i\}$ when there is a value or mark, such as the event’s magnitude or seismic moment, associated with each event time.

A classical point process model of earthquake occurrence is the epidemic-type

aftershock sequence (ETAS) model first proposed by Ogata (Ogata, 1988; Ogata, 1998). The ETAS model and its variants consider earthquake occurrence to be a time-dependent process: earthquakes cause aftershocks, which in turn cause more aftershocks, and so on.

The ETAS model is specified by its conditional intensity, $\lambda(s, t)$, which represents the infinitesimal rate at which earthquakes are expected to occur around time t and location s , given the history H_t of the process up to time t , based on the following simple assumptions (Ogata, 1998): 1) The background seismic activity is generated by a stationary Poisson process with a constant hazard rate; 2) Each shock has a risk of stimulating aftershocks proportional to $e^{\beta M}$, where M is the magnitude of the shock; 3) The hazard rate of aftershocks decreases with time according to the modified Omori law, $\frac{K}{(t+c)^p}$.

It can be viewed as a special case of the linear, self-exciting Hawkes' point process (Hawkes, 1971), where the conditional intensity is of the form

$$\lambda(s, t|H_t) = \mu(s, t) + \sum_{t_i < t} g(s - s_i, t - t_i; M_i), \quad (1.7)$$

where $\mu(s, t)$ is the mean rate of a Poisson-distributed background process that may in general vary with time and space, g is a kernel function that indicates how previous occurrences contribute, depending on their spatial and temporal distances and magnitudes, to the conditional intensity λ at the location and time of interest, and (s_i, t_i, M_i) are the origin times, epicentral locations and moment magnitudes of observed earthquakes. An example form of function g is

$$g(s, t, M) = K(t + c)^{-p} e^{a(M - M_0)} (|s|^2 + d)^{-q}, \quad (1.8)$$

where M_0 is the lower magnitude cutoff for the observed catalog (Ogata, 1998).

The model can be extended in the several ways. “Mutually exciting” processes

can be included in the following form:

$$\lambda_n(s, t|H_t) = \mu_n(s, t) + \sum_j [\sum_{t_i < t} g_{nj}(s - s_i(j), t - t_i(j); M_i)], \quad (1.9)$$

where $n = 1, \dots, M$, $j = 1, \dots, M$ are the different components, and g_{nj} for $n \neq j$ are transfer functions which determine the contribution to the risk of type n events from past events $t_i(j)$ in the j th component. This approach can be used to estimate causal influence of events in one region on the occurrence of events in another (Ogata et al., 1982).

Magnitudes can be incorporated in a marked point process intensity of the form

$$\lambda(s, t, M|H_t) = f(M)[\mu(s, t) + \sum_{t_i < t} h(M_i)g(s - s_i, t - t_i; M_i)], \quad (1.10)$$

where $f(M)$ is the marginal distribution of magnitudes and $h(M)$ is a weighting function which increases rapidly with magnitude.

Another extension is to replace s by a continuous spatial coordinate variable \vec{s} , the location of the fault movement, to yield a space-time-magnitude intensity of the form

$$\lambda(\vec{s}, t, M|H_t) = f(M)[\mu(\vec{s}, t) + \sum_{t_i < t} h(M_i)g(\vec{s} - \vec{s}_i, t - t_i; M_i)], \quad (1.11)$$

where $\mu(\vec{s}, t)$ is a spatially distributed immigration or background noise term, and $g(t, \vec{s})$ is a space-time kernel (Musmeci and Vere-Jones, 1992). Additional variables such as the orientation of the fault movement can be incorporated (Kagan and Jackson, 1994).

The parameters in the ETAS model and its variants can be estimated by maximizing the log-likelihood,

$$\sum_{i=1}^N \log[\lambda(s_i, t_i)] - \int \int_S \lambda(s, t) ds dt. \quad (1.12)$$

This optimization problem can be solved using any of the standard routines, such as the quasi-Newton methods. Under some general conditions, the maximum likelihood estimator (MLE) is asymptotically unbiased, consistent, asymptotically normal and asymptotically efficient (Ogata, 1978).

1.1.3 Point process models in astronomy: gamma-ray bursts

Point processes are also good mathematical models for many astronomical observations. A particularly fascinating subject is the mystery of gamma-ray bursts (GRBs). GRBs are momentary flashes of high-frequency radiation, lasting from a few milliseconds to a few thousand seconds, associated with explosive events in distant galaxies (Klebesadel et al., 1973).

In 1963, suspecting the Soviet Union was conducting covert nuclear weapon tests on the dark side of the moon, the U.S. Air Force launched a series satellites, Vela (*She watches over* in Spanish), to monitor compliance with the 1963 Partial Test Ban Treaty. On July 2, 1967, a flash of gamma-rays was recorded by two of the satellites. The resulting data was noticed and analyzed in 1969 and the discovery of GRBs was published by the gamma-ray astronomy group at the Los Alamos National Laboratory in 1973 (Klebesadel et al., 1973; Bonnell and Klebesadel, 1996).

GRB has been called “the greatest mystery in modern astrophysics” by many astronomers because it is the only phenomenon known to exist, well-studied observationally, but the locations of and the physical mechanism responsible for which remain unclear, in spite of the appearance of over 2000 published observational and theoretical papers (Scargle and Babu, 2003). Only in the past 5 years more coherent hypotheses started to emerge, but new observations of unexpected GRBs continue to challenge current standard theories about the systems that produce these explosions (Ackermann et al., 2014; Maselli et al., 2014).

GRBs have two paradoxical observational characteristics: observed homogeneity

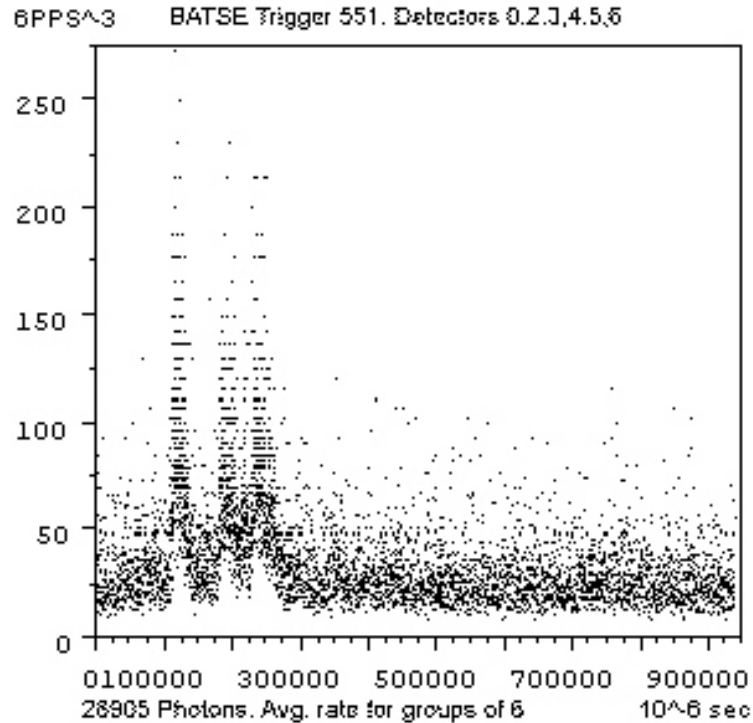


Figure 1-1: Photon counts of a particular GRB event, *Trigger 0551*, as a function of time. (Reproduced from Dr. Jessie Henshaw’s personal archive: <http://www.synapse9.com/>.)

in the Universe with extreme heterogeneity of style of variability. On the one hand, currently orbiting high-energy telescopes detect on average approximately one GRB per day, from random locations on the sky (there are at least as many galaxies in the sky as neurons in human brain). The completely uniform distribution of GRBs in the sky has been a pivotal and surprising observational result that ruled out a whole class of models (Meegan et al., 1992). On the other hand, each GRB exhibits different time history: a single sharp spike, a period of rising intensity followed by a period of decaying intensity in radiation, or a random array of overlapping pulses.

One of the most fundamental data analysis problems in astronomy is the representation of the time variability of a source, or light curve, based on data comprising a series of brightness observations spaced in time. In the context of GRBs, it means

describing the time variation of the brightness of the source of a GRB event, using the arrivals of individual photons.

Figure 1.1 plots gamma-ray intensity for a particular GRB event, *Trigger 0551*, showing three gamma ray burst events lasting about .03 seconds each, a total of 28904 photons (Henshaw, 1998). The plot shows photon counts, summed over four different energy channels (the raw telescope data includes, for each photon, a tag indicating which of four ranges of energy the photon occupied), as a function of time. Each point in the graph represents the average rate based on every 6th photon.

One model that has been adopted to estimate the light curve of a GRB event is called the piecewise constant Poisson model (Scargle, 1998). The idea is to construct an explicit, piecewise constant model for the observed photon data, and marginalize all of the model parameters except for K , the parameter specifying the number of pieces, or blocks. Then, to find the most likely piecewise constant Poisson rate model of the data is simply to find the maximum of this function of K .

The model is a set of Bayesian blocks partitioning T : $M = \{K, n, \lambda\}$. $n = (n_1, \dots, n_K)$ is the K -dimensional array of change-point locations, where $\min(t_n) \leq n_i \leq \max(t_n)$. $\lambda = (\lambda_0, \dots, \lambda_K)$ is the $(K + 1)$ -dimensional array of Poisson rates, either of photons or of the GRB event per unit time, over the corresponding partition segments.

If M_K is the set of all such piecewise continuous models with K change-points without specifying the change-point locations or rates between change-points, then the Bayes factor that determines the status of model M_K relative to any other model is defined as follows:

$$p(D|M_K) = \int p(D|n, \lambda, M_K)p(n, \lambda|M_K)dn d\lambda. \quad (1.13)$$

Bayesian blocks (Scargle, 1998; Dobigeon et al., 2007; Scargle et al., 2013) is a type

of change-point models that detect bursts, reveal pulse shapes, and generally characterize intensity variations in the photon counting data. More general multi-scale methods have since been developed where the underlying intensity function potentially has structure at multiple time- or size-scales (Babu and Feigelson, 1992; Young et al., 1995; Kolaczyk, 1999; Kolaczyk and Nowak, 1999; Kolaczyk and Dixon, 2000; Kolaczyk and Nowak, 2005). Together these methods have brought an interesting non-parametric perspective into the estimation of the intensity function $\lambda(t)$ that is fundamental to the analysis of point process data.

What’s more, these methods focus on estimating a binned version of the Poisson intensity function $\lambda(t)$, rather than the continuous-time conditional intensity function $\lambda(t|H_t)$ discussed in other subsections of this chapter. Two common conceptual schemes in astronomy underlie this choice of modeling approach (Scargle and Babu, 2003). First, interaction between photons traveling through the far reaches of space in the faint beams of radiation is extremely rare, which satisfies the Poisson assumption. Second, this discontinuous form of $\lambda(t)$ provides useful approximations of locations, strengths, and width of the localized pulses, whose overlap is neglectable and whose superposition consists the light curve.

1.1.4 Point process models in neuroscience: spikes

Neuroscience is an additional field of science that produces point process data—firing times of electrical impulses in nerve cells, called spike trains, that interact and form networks. Brillinger published some of the seminal statistical papers on the data analysis of nerve cell spike train data (Brillinger, 1975; Brillinger et al., 1976; Brillinger, 1988b; Brillinger, 1988a; Brillinger, 1992; Brillinger and Villa, 1997).

Two essential questions Brillinger proposed (Brillinger, 1992) are

1. Can an analytic model incorporating the basic features of neuron behavior be developed and fit?

2. Given the firing times of a network of neurons, can one infer their causal connections?

Point process models have been used successfully to address both questions.

For example, point process models can be used to decode neural firing in rat hippocampus during a spatial navigation task (Frank et al., 2001; Huang et al., 2009) and in primate motor cortex during arm reaching tasks (Truccolo et al., 2005). In these cases, the conditional intensity function is used to specify the multiple factors that simultaneously affect spiking activity and to assess the effects and relative importance of each factor. The conditional intensity of an individual neuron can be fitted under the generalized linear model (GLM) framework (Truccolo et al., 2005):

$$\lambda(t_k|N_{1:k}, \theta) = \exp\left\{\gamma_0 + \sum_{n=1}^Q \gamma_n \Delta N_{k-n}\right\} \quad (1.14)$$

where Q is the order of the autoregressive process, γ_n represents the autoregressive coefficients, and γ_0 is the background firing rate. This model can be estimated using likelihood methods and its goodness-of-fit can be diagnosed by using the time rescaling theorem (Brown et al., 2002) in conjunction with Kolmogorov-Smirnov (K-S) test.

By extending the conditional intensity to the multivariate case, we can assess the functional connectivity, or coupling strength, among an ensemble of neurons. For example, Gerhard et al. (2013) built a small network model for three neurons (X, Y, Z) in the crab stomatogastric nervous system of the form:

$$\lambda_Y(t_i) = \sum_{j=0}^{M_s} \vec{\gamma}_{Y,j} \Delta N_Y(i-j) + \sum_{j=0}^{M_c} \vec{\gamma}_{X,j} \Delta N_X(i-j) + \sum_{j=0}^{M_c} \vec{\gamma}_{Z,j} \Delta N_Z(i-j), \quad (1.15)$$

where M is the maximal time lag to consider for coupling, and $\vec{\gamma}_{\cdot,j}$ are the model coefficients for the interaction with spiking activity from neuron \cdot .

The directed coupling strength between two neurons is then defined as

$$CS(X \rightarrow Y) = \left| \int_0^{\infty} \sum_j^{m_c} \gamma_{X,j} D_j(\tau) d\tau \right|, \quad (1.16)$$

where D_j , $j = 1, \dots, m_c$ is some basis function.

This approach, based on spike train data only, for the first time, has been shown to successfully identify relative connection strengths that match the known physiology of the pyloric circuit of the stomatogastric ganglion of the crab even though synaptic transmission in the pyloric circuit is only partly mediated by spikes (Gerhard et al., 2013).

1.2 State-space Models

In many cases, an underlying physical process is more or less continuous, but is only revealed to the observer through discrete events. The discreteness is due to the nature of the observation, and is not intrinsic to the physical process. For example in astronomy, light curves of variable stars are obtained by counting the photons as they arrive at the focal plane of the telescope.

The separation of intrinsic point processes from observation point process leads to the introduction of another class of statistical models: State-space models, which we will now discuss with details illustrated from applications in neuroscience.

1.2.1 Theoretical basics

A common interest of scientists in many fields is to understand the relationship between the dynamics of a physical system and the occurrences of localized events within a physical system. Seismologists study the connection between mechanical vibrations of the Earth and the occurrences of earthquakes so that future earthquakes can be better predicted. Astrophysicists study the oscillating brightness of celestial region

and the emission of photons to understand objects in our Universe. In particular, dynamical analyses of electro-physiological impulses in neural systems have been of increasing interest recently.

Such relationship between the dynamics of a physical system and discrete observations can often be formulated within the framework of state-space models with point process observations. The basic idea is that the dynamics of the physical system is driven by the dynamics of some stochastic state variable \vec{z}_k , and the discrete events ΔN_k we observe in an interval $(t_{k-1}, t_k]$ are a noisy, subsampled version of the state variable.

To illustrate how this framework works, let us consider a problem in neuroprosthetics: Decoding the desired position of a person’s arm by analyzing neural spiking activity from primary motor cortex. In this setting, the state variable \vec{z}_k corresponds to the desired position of the person’s arm at time t_k ; the observation at the current time interval, ΔN_k , and the sequence of previous observations H_k are neural spiking data. An encoding analysis would first describe how position is represented in the neural spiking activity using the conditional probability $p(\Delta N_k | \vec{z}_k)$. A subsequent decoding analysis would make inferences about the movement trajectory through time t_k given all of the observed data up to time t_k using the conditional probability $p(\vec{z}_k | \Delta N_k, H_k)$. Then the estimated desired arm position at any given time t_k can be some function of the conditional probability $p(\vec{z}_k | \Delta N_k, H_k)$, such as its expectation or maximum.

1.2.2 State-space models in the analysis of dynamic neural spiking data

State-space modeling has been successfully applied in a wide variety of neural applications. A detailed review on the subject can be found in (Paninski et al., 2010).

The state-space framework has been applied to perform optimal decoding of state variables that have concrete and real explanations, such as the position of the subject’s

hand in the design of motor neuro-prosthetic devices (Donoghue, 2002) where the state-space models are used to decode from multiunit spiking activity from motor cortex (Truccolo et al., 2005; Wu et al., 2006; Srinivasan et al., 2006; Wu et al., 2009) and parietal cortex (Kemere et al., 2008) and the position of a rat traversing a maze given multiple hippocampal spike train (Brown et al., 1998; Zhang et al., 1998; Eden et al., 2004; Huang et al., 2009).

More recently, this state-space framework have been applied successfully to estimate more abstract neural measures from spike train data. For example, state-space models are used to decode nonstationary neuron tuning properties (Brown et al., 2001a; Frank et al., 2002; Eden et al., 2004; Czanner et al., 2008), where the state variable represents a parameter vector which determines the neuron’s stimulus-response function. Biophysical neural models can be estimated from spikes using state-space models, where the state variable represents a parameter vector, which contains, for example, the membrane voltage and various voltage-gated ion channels, and characterizes a biophysically realistic model of the neuron (Meng et al., 2011; Meng et al., 2014). State-space models are also used to track the probability of learning during behavioral experiments (Smith and Brown, 2003; Smith et al., 2004; Smith et al., 2005; Suzuki and Brown, 2005; Prerau et al., 2009), where the state variable represents the subject’s certainty about the behavioral task. Lastly, by combining multiple state processes, the state-space models can be used to integrate information about the sleep onset process, the gradual transition between wakefulness and sleep (Prerau et al., 2014).

1.3 Goals of the thesis

Despite the success of state-space models in applications, a few statistical challenges remain. First, the classical framework often requires the dynamics of the physical

system to be estimated to be low-dimensional and directly observable during the encoding step. This hinders our ability to understand abstract, higher-order processes that are unobservable, such as various cognitive processes as studied in the human brain. Second, as recording technologies advance, the observed data are now growing in dimension and variability, making it difficult to directly relate to relevant covariates simultaneously. The classical framework has yet to incorporate new methods related to dimensionality reduction. Third, the relationship between the dynamics of the physical system and the observation is often of an unknown, non-linear function that is also time-varying. Lastly, specifically in neuroscience, closed-loop experiments that have emerged in recent years seek to characterize causality between neural activity and behavior. This requires decoding methods that can make decision in real-time (for example, to stimulate the neurons or not) based on various sources of information present in the data.

In this thesis, we propose several new methodological developments that advance state-space point process methods to address these challenges. We illustrate the applications of these methods using data from neurophysiological investigations. A schematic outline of the four sections is shown in the figure below.

In Chapter 2, we apply the state-space framework to a novel problem of tracking rhythmic dynamics in the spiking of neurons in the subthalamic nucleus of Parkinson’s patients. In this case, the state variable \vec{z}_k represents a parameter vector which determines how a neuron’s history spiking activity at different time lags modulate current spiking.

In Chapter 3, 4, and 5, we extend the state-space framework to accommodate marked data, to allow for real-time decision making, and to take into account the need for dimensionality reduction, respectively. The statistical intuition comes from the fact that the framework of state-space models can be broken down into two

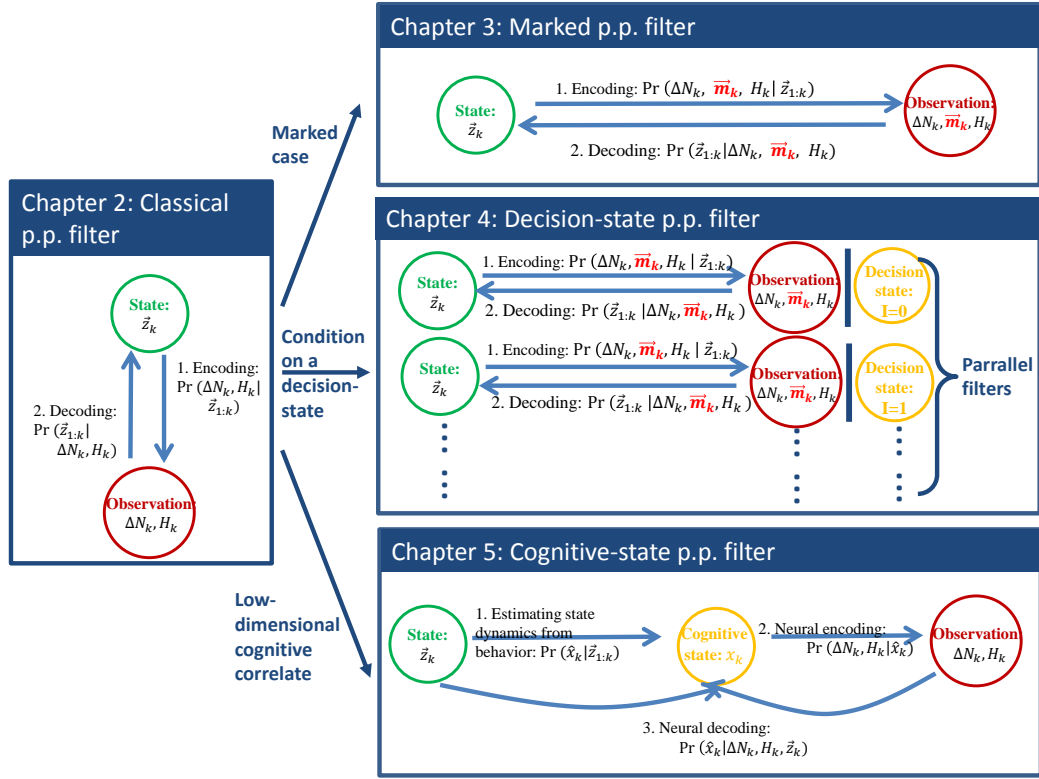


Figure 1.2: Schematic outline of the thesis.

main components: the state model and the observation model. The state model describes the time evolution of the state variable. The observation model, $p(\Delta N_k | \vec{z}_k)$, characterizes the distribution of the observed discrete events as a function of the state.

In Chapter 3, we generalize the observation model to marked point processes $\{\Delta N_k, \vec{m}_k\}$ where \vec{m}_k represents a collection of mark vectors that characterize the event arrival. In the case of tetrode recordings, for example, the mark could be a length four vector of the maximum amplitudes of the spike waveforms on each of the four electrodes at every spike time. In Chapter 4, we extend the state variable from a single continuous state z_k to a composite vector $\{z_k, I\}$ where z_k remains an underlying dynamic, continuous state process and I is a discrete, fixed decision state variable. We make use of the underlying state process in a content-specific way by selectively including different sources of information. We break down the decoder

into components and investigate the contributions of each component to decode the discrete decision state. In Chapter 5, we introduce an intermediate state variable x_k to handle the growing dimensions of both the dynamics of the physical systems and the observed data. We present a framework relating neural activity to more general and abstract variables that influence multiple aspects of behavior and cognitive function by introducing a cognitive state.

Chapter 2

A Point Process Approach to Identifying and Tracking Transitions in Neural Spiking Dynamics in the Subthalamic Nucleus of Parkinson's Patients

2.1 Introduction

Over recent years, a great appreciation of brain rhythms has emerged in the field of neurological disease. Gradual shifts in lower frequency activity in alpha (7–12 Hz) and beta (12–30 Hz) band oscillations are a well-known hallmark of Alzheimer's disease (Berendse et al., 2000; Coben et al., 1983; Huang et al., 2000; Osipova et al., 2005; Penttila et al., 1985; Schreiter-Gasser et al., 1993). Low frequency oscillations (3–10 Hz) may contribute to Parkinsonian tremor whereas enhanced beta band oscillations in the basal ganglia and cortex correlate to bradykinesia and rigidity in Parkinson's disease (Brown et al., 2001b; Levy et al., 2002) and may result from amplification of normal striatal network dynamics (McCarthy et al., 2011). Patients with schizophrenia have shown a reduced amplitude of gamma (30–200 Hz) and theta (4–7 Hz) oscillations in frontal regions (Haenschel et al., 2009; Schmiedt et al., 2005; Cho et al., 2006; Winterer et al., 2000; Uhlhaas and Singer, 2010).

To study such rhythmic dynamics in patients with various neurological disease, a number of statistical data analysis techniques are common, many of which focus on spectral estimation methods such as the periodogram or multi-taper estimation,

or on autoregressive modeling. However, this wide array of signal processing tools are most appropriately applied to continuous-valued recordings of electrophysiological activity such as magnetoencephalography (MEG), scalp electroencephalography (EEG), surface electromyography (EMG), and local field potentials (LFP) (Brown et al., 2001b; Levy et al., 2002; Salenius et al., 2002; Pfurtscheller, 1981; Marsden et al., 2001; Foffani et al., 2005), rather than spike trains, which take on discrete values in time. Similarly, deterministic computational modeling methods for neural electrophysiological rhythms are often formulated in terms of differential equations or difference equations, which are more appropriate for continuous-valued signals. Many standard analysis methods available for neural spike train data assume that the firing structure is static in time, limiting their ability to identify and track dynamic transitions. Hence, there is a considerable need for dynamic statistical methods explicitly designed to analyze neural spike train data (Brown et al., 2004; Brown, 2007).

The analysis of the dynamic properties of spike data presents a number of critical statistical challenges. First, as opposed to continuous-valued measurements, statistical estimation and inference procedures for neural spike train data are most appropriately developed based on the theory of point processes. Inferences based on standard methods developed for continuous valued signals can lead to a reduction in the statistical power available in the data or to inappropriate or incorrect conclusions about associations in the data. For example, spectral analysis methods applied to point process data present particular issues not present for continuous valued signals. An idealized spike contains power at all frequencies (Brillinger, 1975), making it difficult to interpret the structure of stochastic spiking properties from spectral estimates in a narrow frequency band. Often, analyses of spike train data using spectral methods assume that the data arises from a Poisson process, restricting the types of dynamics that can be captured, and often resulting in smaller than expected power

in the low frequencies (Bair et al., 1994; Franklin and Bair, 1995). Such spectral distortion makes it difficult to identify rhythms in the low frequencies (Rivlin-Etzion et al., 2006). Additionally, neural spike train data exhibits a wide variety of history dependent behaviors such as refractoriness, bursting, and intrinsic rhythms. It has been shown that the autocorrelation function of the high-frequency discharge in globus pallidus cells contains many small peaks that do not necessarily reflect bursting activity, and could be inappropriately interpreted as rhythmic spiking (Bar-Gad et al., 2001). However, the majority of the continuous-valued data analysis methods, such as tests based on spectral estimators, are performed under an assumption of stationarity. Such assumptions are rarely true of history-dependent point processes. The violation of stationarity by history-dependent neural spike train data can lead to false detections of changes in the rhythmic dynamics or reduced power in classic testing methods.

Second, the probability of a neuron firing a spike in any time interval is often influenced by many factors occurring simultaneously. These can include influences from external biological or behavioral signals, from the neuron's past activity, or from the activity of other interacting neurons. Many data analysis techniques for spike train data, which focus on the influence of a single factor without regard for others, can lead to incorrect inferences. For example, the coherence between a spike train and an LFP is influenced by the average firing intensity (Lepage et al., 2011). An increase in neural firing rate can also have the effect of increasing the maximum oscillatory frequency that can be transmitted by a neural spike train, as observed in LFP recordings from the subthalamic nucleus (STN) of Parkinson's patients (Levy et al., 2002). Therefore, spike trains in STN can form robust rhythms by modulating their firing intensity by small amounts. This makes detecting and tracking neural dynamics difficult.

To address these issues, we propose a state-space point process approach for identifying and tracking changes in rhythmic neural dynamics in spike train data. The approach uses history dependent point process models, as rhythmic spiking dynamics are closely related to the history-dependent firing activity. In terms of the statistical point process model, rhythmic neural activity related to physiological oscillations is reflected as a structured history dependence between past and current spiking. This is true for continuous valued signals such as LFPs, where spectral characteristics are reflected in the autocovariance of past and current signal values. Whereas the autocovariance function accounts for history dependence in the signal at a single temporal lag, a statistical model might account for such dependence across multiple lags. Similarly the point process model that we develop here characterizes the rhythmic spiking activity in terms of history dependence across a range of temporal lags. Here, we fit rhythmic spiking in Parkinsonian STN neurons to a point process model that captures the modulation of the firing probability as a function of previous spiking across multiple lags. We uncover a consistent dynamic pattern of modulation that characterizes the structure of and suggests potential mechanisms for beta frequency oscillations in the spiking data.

The proposed approach is as follows: First, we construct a point process model with two sets of parameters defining the influence of past history in two distinct physiological states. We estimate these parameters as well as the state transition times that maximize the likelihood of the observed data. We then develop a maximum likelihood ratio test to determine whether such a model provides a significant improvement over one that has only a single set of history parameters and no dynamical transitions. Then, if there exists a statistically significant difference between the two models, we apply a point process filtering and smoothing algorithm to track the temporal dynamics of the transition through time. This two-step approach is

illustrated with simulated data and with real spike train data recorded in the STN of Parkinson’s patients performing a hand movement task.

2.2 Methods

Experimental data were obtained from single neurons in the subthalamic nucleus (STN) of patients undergoing surgery for the treatment of Parkinson’s disease (PD). Details of the patient selection, basic recording protocols and behavioral paradigm, and previous analyses of this experimental data are discussed in Amirnovin et al. (Amirnovin et al., 2004). The patients viewed a computer monitor and used a contralaterally mounted joystick to perform a hand movement task. At the beginning of each trial, a small fixation point appeared in the center of the monitor. After a 200 ms delay, four small grey targets appeared circularly around the fixation point. After another 300 ms delay, a randomly selected target turned green and a small cursor appeared in the center of the monitor. The patient then used the joystick to guide the cursor from the center of the monitor toward the green target. The patient typically reached the target within 1 s of the green cue. Between trials there was an interval of 1 s and the patients were required to return the joystick to the center position. Patients typically performed 12–24 correct trials in each direction at a given site and 1–4 sites were recorded simultaneously. The initial recording run lasted about 3–5 min for each trial, but in this study we only used recordings of 1.6 s before and after the start of hand movement, which was determined by an initial 1 degree deflection of the joystick (from a possible range of 18 degrees). Spikes were sorted using a template-matching algorithm refined by principal component analysis and cluster cutting. Seventy percent of cells either had clear single units or represented multiunit recordings that were dominated by a single cell. The remaining 30% represented multiunit recordings that could not be refined further. Because we expected to see

history dependent structures reflected in both single unit and multiunit recordings, we included both of the groups in our analyses.

While this movement task included trials in four different directions, our interest here is in the transitions between dynamic rhythmic states, which we posit to be similar for all trials. We therefore pooled all of the trials across all directions, and examined the firing properties of each neuron relative to movement initiation.

For each neuron, we constructed raster plots, peri-stimulus histograms (PSTH), autocorrelation function plots, and spectrograms (periodogram estimators over 100 ms sliding windows) for all trials. We proceeded to examine 22 neurons that exhibited some degree of rhythmic spiking dynamics based on preliminary results from these standard feature detection methods.

2.2.1 Detecting changes in history-dependent firing patterns

In this section, we introduce the statistical methods to assess whether there has been a significant change in rhythmic spiking dynamics. First, we formulated a point process model with history dependent firing properties, under the assumption that a sudden change in dynamic states occurs at some unknown point in time during the observed spike train time series, and then transitions back to the original dynamic state back at some later time point. We call this a two-state point process model. We used maximum likelihood methods to estimate model parameters related to a background firing rate, to the history dependent firing properties, and to the time points of the transitions. We then performed a maximum likelihood ratio test to determine whether there was a significant difference between this two-state model and a single state model that assumes no change occurs in the neurons firing properties.

In order to reduce the dimension of the statistical model, both to explain the data parsimoniously and to limit the chance of model overfitting, we developed a statistical model that related the instantaneous probability of a spike at any time point to the

past spiking history using a set of cardinal spline basis functions. Splines are locally defined third-order polynomial functions that flexibly approximate arbitrary smooth functions using a small number of basis functions (Ramsay and Silverman, 2010). Because our study focused on rhythmic activity, we chose a spline basis set that had been previously used to successfully characterize the desired rhythms in neural spiking activity (Eden et al., 2012). In this previous study, we constructed a point process model that assumed that the history dependent modulation was constant over known time intervals for each trial. This previous model was developed to visualize specific features of the history dependent structure, but not to determine whether the modulation parameters underwent statistically significant changes during movement or to fully characterize the time course of these changes. In this study, we extended the previous model in two ways: first, we estimate the transitions between the non-movement and movement state and use a nested set of models to determine whether the data shows statistically significant changes in the modulation parameters; second, we develop a state space model and a point process filtering and smoothing algorithm to track the dynamic properties of changes in the modulation parameters through time.

A statistical model for point process data is defined by a conditional intensity function $\lambda(t|H_t)$, which characterizes the instantaneous spiking probability at any time, t , as a function of t , the past history of spiking H_t , and any other factors that may influence spiking activity (Daley and Vere-Jones, 2003). In this case, the conditional intensity model was defined as follows:

$$\log \lambda(t|H_t) = \beta_0 + \gamma_0 \cdot I_{\text{state}}(t) + \sum_{i=1}^p \beta_i \cdot G_i(H_t) + \sum_{i=1}^p \gamma_i \cdot G_i(H_t) \cdot I_{\text{state}}(t), \quad (2.1)$$

where H_t is the neuron's spiking history going back 100 ms, β_0 relates to a background intensity of spiking, $I_{\text{state}}(t)$ is the binary indicator function that is equal to 0 during

an initial non-movement state, becomes equal to 1 when the neuron transitions to a movement planning and execution state, and returns to 0 when the neuron transitions back to the non-movement state. β_i is the parameter multiplying the spline basis functions at the i -th control point. When exponentiated, e^{β_i} can be interpreted as modulation of the intensity due to a previous spike at the specified lag during the non-movement period. γ_i represents the baseline and history modulation coefficients for this indicator function, p is the number of spline control points, and G_i are a set of spline basis functions at multiple selected temporal lags. Here we assumed the history-dependent spiking activity before and after the movement planning and execution state were the same, thus were expressed by the same set of model parameters, while the rhythmic activity during the movement state is fitted with a different set of parameters. Moreover, p was set to 10 based on previous spline-based models developed with this data (Eden et al., 2012).

If $I_{\text{state}}(t)$ is known, this formula describes a generalized linear model (GLM) for the spike train data (Truccolo et al., 2005). Such GLMs have a number of nice properties, including convexity of the likelihood surface and asymptotic normality of the parameter estimates. However, there are two additional unknown parameters, the transition times into and back from the movement state. We computed maximum likelihood estimators for these parameters numerically, by computing all model parameter estimates over a grid of possible values for these transition times. The rows of the grid contained the possible transition start times, which spanned a minimum value of 700 ms before movement onset to a maximum value of 500 ms before movement onset, with stepwise increments of 50 ms. The columns of the grid contained the possible transition end times, which spanned a minimum value of 800 ms after movement onset to a maximum value of 1000 ms after movement onset, with stepwise increments of 50 ms. At each point in the grid, we calculated the log-likelihood of the

corresponding transition start and end times, maximized over the remaining baseline rate and modulation parameters. This provides a log-likelihood surface for the transition times. We identified the start and end transition times that maximized this log-likelihood surface and used those as point estimates for the fit model. Confidence intervals were computed from the estimated Fisher information for the baseline and history modulation parameters and by computing a local quadratic approximation of the log-likelihood surface for the transition times (Le Cam, 1960; van der Vaart, 1998). Using these maximum likelihood estimates to identify the movement period, we compared the two different sets of parameter estimates for history dependent modulation of firing intensity in the movement and non-movement states. In particular, we focused on changes in the background firing rate parameter, as well as changes of modulation at time lags 2 ms, 30 ms and 50 ms, which may relate to bursting behavior and to rhythmic spiking in the beta frequency band. We also computed the ratios of those parameter estimates between the two states.

To assess whether the two states are significantly different, we tested the hypothesis that the neurons go through a different dynamic state during movement planning and execution than before and after movement against the null hypothesis that there is no transition of dynamic states throughout the experiment. The hypotheses can be formulated as follows.

$$H_0 : \gamma_i = 0 \text{ for all } i \quad \text{v.s.} \quad H_A : \text{For at least some } i, \gamma_i \neq 0. \quad (2.2)$$

Because the null model is nested under the proposed model, we performed a maximum likelihood ratio test where the test statistic is defined as follows:

$$\Lambda = -2 \log \frac{L(\lambda(t|H_t)|H_0)}{L(\lambda(t|H_t)|H_A)} = -2 \log L(\lambda(t|H_t)|H_0) + 2 \log L(\lambda(t|H_t)|H_A) \quad (2.3)$$

where $\log L(\lambda(t|H_t)|H_0)$ is the log-likelihood for null model and $\log L(\lambda(t|H_t)|H_A)$ is

the log-likelihood for alternative model. Here the test statistic follows an asymptotic chi-square distribution with 11 degrees of freedom (Le Cam, 1960; van der Vaart, 1998).

To assess goodness-of-fit of the two-state model, we constructed Kolmogorov-Smirnov (K-S) plots of time-rescaled inter-spike intervals (ISIs) (Brown et al., 2002). The time-rescaling theorem produces a set of rescaled ISIs that are independent with an exponential distribution with mean 1 if the proposed model accurately describes the structure in the observed spiking activity. To construct the K-S plot, we plot the empirical cumulative distribution of the rescaled ISIs against the theoretical cumulative distribution of the Exponential(1) distribution. The better the quality of the model fit, the closer the K-S plot should be to a 45 degree line (Johnson and Kotz, 1980).

In addition to the observed data, we also applied the two-state model to simulated data. Two hundred spike trains were simulated discretely from a Poisson process with unit rate and then the spike times were rescaled according to a model conditional intensity function (Brown et al., 2002). The parameters of this simulation model were chosen to reflect features observed in the real data. We then estimated the model parameters from the simulated data, and compared these estimates to the known model parameters used to generate the data.

Additionally, in order to study how large a change in history dependent modulation can be detected by the maximum likelihood ratio test procedure in Equations 2.2 and 2.3, we conducted a power analysis using the simulated spike train data. We sampled a variety of different values of modulation parameters related to the beta rhythm (γ_i corresponding to lag 30 ms and 50 ms) and in each case generated 40 simulations of 200 trials of spike trains. We then calculated the proportion of simulations for which we were able to successfully reject the null hypothesis.

2.2.2 State-space smoothing algorithm to track the change through time

The two-state model described above assumes that the transitions between dynamic states occur suddenly and at identical times across all trials. To allow for more gradual and trial dependent transitions between dynamic states, we developed an alternate, adaptive model, where the modulation parameters, β_i , are allowed to change smoothly from one time point to the next. That is, we assumed that the parameters responsible for the rhythmic spiking dynamics are time-varying and we estimated the parameters at each point in time. The conditional intensity function for this adaptive model is

$$\log \lambda(t|H_t) = \beta_{0,t} + \sum_{i=1}^p \beta_{i,t} \cdot G_i(H_t). \quad (2.4)$$

Here the modulation parameters $\beta_{i,t}$ are time dependent. We estimate the set of parameters at each time point on a millisecond timescale using a stochastic state point process filtering and smoothing algorithm (Eden et al., 2004). The algorithm recursively estimates the parameters based on the previous estimates and the instantaneous data at each time point. It can be broken down into an initial filtering component,

$$\theta_{t|t-1} = \theta_{t-1|t-1} \quad (2.5)$$

$$W_{t|t-1} = W_{t-1|t-1} + \Sigma \quad (2.6)$$

$$W_{t|t} = \left[W_{t|t-1}^{-1} + \frac{\partial^2 \log \lambda}{\partial \theta \partial \theta^T} (\Delta N_t - \lambda_t \Delta t) + \left(\frac{\partial \log \lambda}{\partial \theta} \right)^T \lambda_t \Delta t \left(\frac{\partial \log \lambda}{\partial \theta} \right) \right]^{-1} \quad (2.7)$$

$$\theta_{t|t} = \theta_{t|t-1} + W_{t|t} \cdot \frac{\partial \log \lambda}{\partial \theta} (\Delta N_t - \lambda_t \Delta t) \quad (2.8)$$

following by a smoothing component,

$$\theta_{t|T} = \theta_{t|t} + W_{t|t} W_{t+1|t}^{-1} (\theta_{t+1|T} - \theta_{t+1|t}) \quad (2.9)$$

$$W_{t|T} = W_{t|t} + W_{t|t} W_{t+1|t}^{-1} (W_{t+1|T} - W_{t+1|t}) W_{t+1|t}^{-1} W_{t|t} \quad (2.10)$$

where $\theta_{t_1|t_2}$ represents the estimated mean of the vector of β_i parameters at time t_1 using all of the data up to time t_2 , and $W_{t_1|t_2}$ represents the estimated covariance matrix of these parameter estimates at time t_1 using all of the data up to time t_2 . Σ represents a one-step covariance matrix that defines the possible evolution of the modulation parameter vector from one time step to the next. The first and second partial derivatives of $\log \lambda$ with respect to the parameter vector θ are all evaluated at $\theta_{t|t-1}$. We initialized the algorithm by setting $\theta_{0|0}$ and $W_{0|0}$ to be our estimates and covariance matrix from the two-state model.

To evaluate the adaptive changes, we plotted the estimates for the background firing rate and the estimates for the modulations due to past spiking between 11 ms to 100 ms in the past against time. In order to highlight statistically significant features of these estimates, we compute point-wise significance levels for the estimates of the modulation at each lag and time point. These point-wise bounds are computed by pre- and post-multiplying the posterior covariance estimates, $W_{t|T}$, with the matrix of spline basis functions, and taking the diagonal elements. Note, since these are point-wise estimates, a number of these will achieve significance by chance alone. While one could develop more stringent criteria to account for this multiplicity, here our goal is simply to highlight large statistically significant regions where the modulation properties have undergone dynamic changes. To quantify the uncertainty of the estimated parameter transitions, we also plotted the standard deviations of estimates of background firing rate, and select history-dependent parameters at time lags of 2 ms, 30 ms and 50 ms against time.

As with the two-state model, we used the resulting model estimates to construct K-S plots of the time-rescaled ISIs for the adaptive model, both for the real data, and for simulated data from 200 simulated spike trains to illustrate that the parameters of the adaptive model could be successfully recovered.

2.3 Results

2.3.1 Preliminary analysis

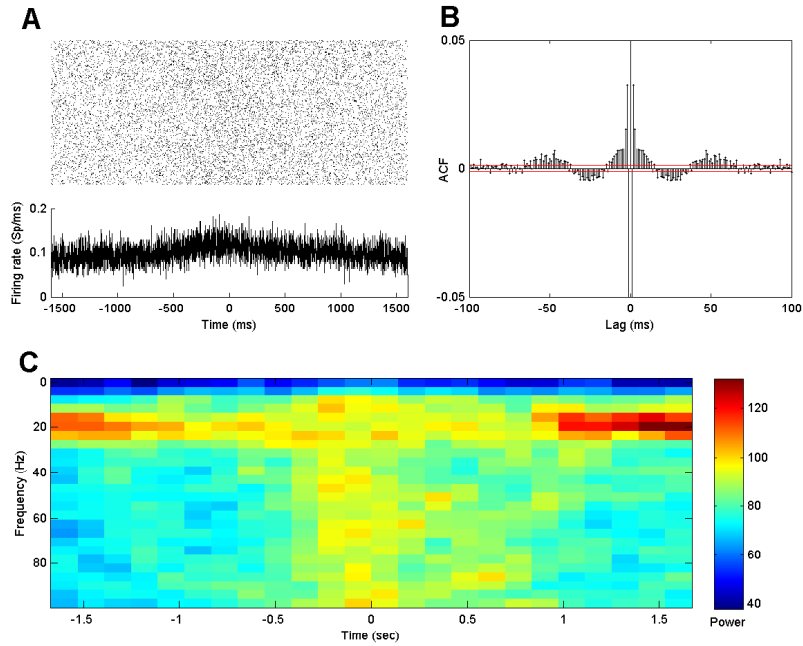


Figure 2.1: (A) Raster and peri-stimulus histogram (PSTH) of 3.2 second recordings from a single neuron before and after onset of hand movement (time 0) for all trials. (B) Trial-averaged autocorrelation functions. The red lines denote pointwise 95% significance bounds. (C) Trial-averaged spectrograms of spike trains as a function of time relative to movement onset on the x-axis and frequency between 0–100 Hz on the y-axis.

Figure 2.1A shows a raster plot of the spiking activity across all trials for an example neuron. Below these rasters is a peri-stimulus time histogram (PSTH) aligned to the onset of the joystick movement. Visual inspection of the PSTH shows that the firing rate increases slightly at about 600 ms before onset of hand movement and then subsides back to its previous value about 600 ms after the onset. Changes in the rhythmic firing properties are not obvious, either in the individual rasters or in the histograms averaged over trials.

Since it is difficult to visualize rhythmic activity in single trials, we examined a variety of standard visualizations, including trial-averaged autocorrelation plots and spectrograms averaged across all trials. Figure 2.1B displays the autocorrelation function (ACF) of the example neuron. The ACF shows a clear negative value at lag 1 ms, followed by a sequence of large positive values extending to 10 ms. The dip at lag 1 ms can be explained by the refractory period of the neuron, and the following increased autocorrelation may correspond to a period during which rapid bursts of spikes tend to appear. More subtle fluctuations in the autocorrelation function are observed for lags between 10 ms and 100 ms, which are associated with spiking rhythms. The autocorrelation is significantly positive for lags between 2–15 ms and 40–60 ms, and is significantly negative for lags between 20–40 ms.

Figure 2.1C shows an estimated spectrogram of the spike train as a function of time on the x-axis and frequency between 0–100 Hz on the y-axis. There is clear beta band activity, illustrated by the increased power between 10–30 Hz, which starts to diminish around 0.7 seconds before the onset of the hand movement, and returns around 0.7 seconds after movement initiation. Power in the gamma band, between 60–100 Hz, exhibits the opposite trend. It intensifies around 0.5 seconds before movement onset and then recedes around 0.9 seconds after movement onset. This high frequency behavior is not limited to a narrow band of frequencies in the Gamma range, but is also present at higher frequencies up to the Nyquist cut-off (500 Hz). This likely reflects, in part, an increase in the expected firing rate during movement. Looking across time, we can visually segment the activity into two seemingly distinct rhythmic states. One state is present during a short planning period before the movement onset and during the movement, where the beta rhythm is diminished and high frequency activity is present. The other state occurs before the movement planning period and returns after movement. Commonly, for spike train data analysis, the spectrogram

is used as a descriptive technique to visualize rhythmic data rather than as the basis for statistical inference. However statistical significance of spectral estimators can be evaluated in a number of ways including using short sliding windows on surrogate data (Hurtado et al., 2004; Hurtado et al., 2005). We do note that inferential procedures based on spectral estimators of spiking data raise a number of issues as discussed in Jarvis and Mitra (Jarvis and Mitra, 2001). Here our goal is instead to develop a model based approach to characterize the influence of past history on the firing intensity directly.

To address this question, we developed a statistical test based on a two-state point-process model, whose application we illustrate both in simulation and to the STN data.

2.3.2 Identifying significant changes in spiking dynamics— simulation study

In order to characterize the properties of the parameter estimation and hypothesis test based on the two-state point process model, we performed a simulation study by generating spike train data from the model with parameter values that capture expected history dependent structure and that undergo different amounts of change between the two dynamic states.

For each parameter set, we simulated 200 spike trains of length 3.2 s using the time-rescaling algorithm described in the methods section. We chose the start and end points of the movement state to be from 400 ms before to 400 ms after movement initiation, and we chose the baseline firing rate to be a constant 50 spikes per second. Both the non-movement and movement states included identical parameters for the first 10 ms of history dependence, which were designed to capture a refractory period and bursting activity typical of these neurons. For the movement state, the remaining history parameters were set to 0, indicating no modulation of intensity based on

previous spiking at these lags. For the non-movement state, the modulation parameter at 30 ms lag was set to $e^{-\alpha}$ and the modulation parameter at 50 ms lag was set to e^{α} , for different values of α , indicating beta frequency modulation of the sort observed previously in the STN (Eden et al., 2012).

From the simulated spike trains, we estimated the parameters for the start and end points of the transition between the states as well as all of the model parameters for the conditional intensity function in Equation 2.1. Figure 2-2A shows estimates and confidence bounds for the transition points between the states for four simulation examples where the parameter α was set to 0.1. This represents a maximum modulation at lags above 30 ms of 10.5% above the baseline firing intensity. In each case, the true transition times of -400 ms and $+400$ ms were accurately captured by the 95% confidence intervals. However, the confidence bounds are quite large, spanning periods from -600 ms to -200 ms for the transition to the movement state, and -50 ms to 750 ms for transitions from the movement state. This suggests that estimation of sharp transitions between dynamic states can be inaccurate, even with a large amount of data.

Figure 2-2B shows a comparison between the true history dependent modulation function of firing intensity and the estimated modulation function with 95% confidence bounds for a single simulated neuron. The cyan dashed line represents true modulation during non-movement state and the magenta dotted line represents true modulation during movement state. The cyan transparent region represents 95% confidence bounds for the estimated modulation during the non-movement state and the magenta transparent region represents 95% confidence bounds for the estimated modulation during the movement state. The estimates tend to follow the true parameter values closely, and as opposed to the estimates of the transition times between states, the confidence intervals for the actual history dynamics are very narrow.

Figure 2.2C shows a K-S plot of the rescaled ISIs for the estimated point process model. The K-S plot follows the 45 degree line closely, with the largest deviation coming from the smallest rescaled ISIs. This could indicate some minor model misfit associated with the refractoriness and bursting. Indeed, when simulated at a temporal rate higher than 1 ms, the deviations at lower rescaled ISIs disappeared (not shown).

Running the maximum likelihood ratio test on the data from this simulation example gives a test statistic of $\Lambda = 123$. If the null hypothesis of no difference between the movement and non-movement states were true, we would expect the value of this test statistic to come from an approximate Chi-square distribution with 11 degrees of freedom. The p-value for the observed test statistic is extremely small (10^{-8}), indicating that there is considerable evidence of a real change in the history dependent spiking dynamics in this data. This is not surprising, given the difference in the estimates and the high degree of confidence in the history dependent parameters between the two states.

Figure 2.2D shows the power of the test procedure as a function of the amount of modulation in the history parameters related to the beta rhythm between the two states (achieved by varying the α parameter). In each case, 40 simulations of 200 trials were used to perform the test. The x-axis indicates the modulation level at the 30 and 50 ms history parameters in the non-movement state and the y-axis indicates the proportion of tests that were able to successfully reject the null hypothesis. It is evident that even for small value of the modulation parameters (2.4%), we are very likely to identify actual changes in rhythmic dynamics, with the amount of data we recorded for this task.

2.3.3 Identifying significant changes in spiking dynamics—STN data

Next, we applied the two-state model to our observed spike train data. Once again, we used maximum likelihood to estimate all of the model parameters and obtain

confidence bounds on those estimates. For the transition times between states, we estimated the maximum likelihood values numerically, and used a local quadratic approximation to the log likelihood to estimate the Fisher information and confidence bounds, as described in the Methods section.

The estimated transition points for the start and end of the movement periods for a set of four representative cell are shown in Figure 2-3A along with 95% confidence bounds. These estimated values and confidence levels are fairly consistent across all of the neurons we analyzed. The confidence intervals for the transition points are all relatively large, spanning about 200 ms for the transition to the movement state and about 300 ms for transitions back. This could arise for multiple reasons including variability in the transition times between trials or a slow rather than sudden change in the firing properties. We will explore the latter possibility in the next sections. However, it is noteworthy that the size of these confidence bounds are not larger than those estimated in the simulation study, where the transitions were instantaneous and consistent across trials. This suggests another possibility—that the state changes associated with altered spiking history dependence are simply difficult to estimate accurately with the amount of data collected in this experiment.

Though the confidence intervals are relatively large for the transition points, the confidence intervals on the model parameters reflecting history dependent dynamics are much smaller and illustrate a clear change between the movement and non-movement states. Figure 2-3B shows the estimated modulation of the firing rate on the y-axis based on a previous spike at any time lag on the x-axis. The estimated modulation during the non-movement state is shown as a cyan dashed line, and the 95% confidence levels are illustrated by the surrounding transparent cyan region. For the movement state, the estimated parameters are given by the dotted magenta line with a 95% confidence level given by the magenta transparent region. Visually, the

non-movement state displays a longer bursting period that ranges from a 2 ms to a 10 ms time lag, a modulation significantly less than 1 for time lags between 10 ms and 40 ms, and a modulation significantly larger than 1 for lags between 40 ms and 60 ms. During the movement state, the estimated modulation exhibits a significant reduction in the parameters above 10 ms related to rhythmic activity. The clear visual separation between the transparent 95% confidence regions between these two states from 10 ms to 60 ms lags suggest that this likely represents a statistically significant change in dynamics.

Figure 2.3C shows an example K-S plot for the model fit of the neuron whose estimated history dependent properties are shown in Figure 2.6B. As with the simulation study, the model and empirical CDFs demonstrate a good overall fit, with some evidence of misfit in the smallest rescaled ISIs. This may suggest some model misfit related to our assumptions of consistent and sudden state changes, or may be related to the difficulty in estimating the transition times with a high level of accuracy.

Figure 2.3D shows the p-values for the significance tests for each of the 22 rhythmic neurons whose firing properties were analyzed. In 21 of these 22 hypothesis tests, we obtained p-values below 0.05 and in 15 of these 22 cells, we obtained p-values below 10^{-8} .

In addition, we analyzed the modulation in the estimated model history parameters between the non-movement and movement states for parameters related to the baseline firing intensity, bursting, and the 30 ms and 50 ms lags related to beta rhythmicity across all 22 STN neurons. Figure 2.4 shows a set of histograms across cells for each parameter, where a value of 1 represents no modulation in the estimated parameter value between non-movement and movement states, a value greater than 1 indicates an increase in that parameter for a particular neuron with movement, and a value less than one indicates a decrease in that parameter with movement.

The upper right panel shows that the baseline firing intensity tends to increase in the majority of cells with movement, with a mean increase of 5.6% from the non-movement baseline firing. The upper right panel shows that modulations 2 ms after a spike related to bursting behavior tend to decrease for all but two cells when entering movement state, with a mean decrease of 10.8% from the non-movement state. The lower panels show changes in the parameters related to rhythmic firing in the beta range related to 30 ms lags (left) and 50 ms lags (right). Most neurons show an increase in the 30 ms lag modulation parameter corresponding to an attenuation of the previously observed inhibition at this lag during movement. The average increase in modulation across all cells is 6.6%. For the parameter related to a 50 ms lag, there are a large fraction of neurons that increase their modulation during movement and a larger percentage that decrease their modulation. The average change in modulation is a decrease of only 0.5%. These results suggest that rhythmic spiking in the beta frequency range is associated with a process of inhibition and excitation, and that changes in the inhibitory component lead to modulation of beta firing more robustly across the neurons we examined. Overall, even though the parameters responsible for rhythmic firing activity represent fairly small modulation in the firing probability at any instant, the two-state model is able to detect these differences accurately.

2.3.4 Tracking changes in spiking dynamics—simulation study

In order to characterize the ability of the state-space framework to track evolving rhythmic spiking dynamics, we generated simulated spike train data that smoothly transitioned from the dynamics of the non-movement and movement states observed when we applied the above two-state model to the data. We simulated 200 spike trains of length 3.2 s using the time-rescaling algorithm described in the methods section. We chose a constant baseline firing rate of 50 spikes per second for all trials. Parameters for the first 10 ms of history dependence were set to constant values

across time, selected to capture the refractory period and bursting activity observed previously in STN neurons (Eden et al., 2012). We simulated changes in rhythmic spiking in the beta frequency range by letting the parameter associated with the 30 ms lag vary quadratically from a starting point of a maximum modulation that represents a 10% decrease in firing probability for a spike 30 ms in the past during non-movement to no modulation for a previous spike 30 ms in the past at the start of movement, and then back to a 10% negative modulation by the end of trial. At the same time, we set the modulation parameter associated with the 50 ms lag to vary quadratically from a 10% increase in the firing probability for a spike 50 ms in the past to no modulation, back to a 10% increase. The remaining history parameters were set to 0, indicating no modulation of intensity based on previous spiking at these lags.

From the simulated spike trains, we used the point process filter and smoother described in Equations 2.4–10 to estimate the trajectory of the state process representing all of the model parameters at each instant in time. Figure 2.5A shows estimates and confidence bounds for the parameter related to the baseline firing rate. The true baseline firing rate, shown as the black line, was held constant at 0.05 spikes/ms. The estimated baseline parameter evolves in time, but hovers at a value close to the true value, with a 95% confidence bound that contains the true value at all times.

Figure 2.5B shows both the true and estimated history dependent modulation of firing intensity as a function of time relative to movement onset on the x-axis and lags between 10-100 ms on the y-axis. The bottom panel shows the true modulation parameter values used to simulate the data. The top panel shows the estimated modulation parameter values at each lag and time point. The middle panel shows only those parameters whose point-wise significance, computed using the estimated posterior covariance matrix, attained a p-value less than $p = 0.05$, with the rest of

the modulation values set to 1. For each of the plots, modulation higher than 1, representing an increase in firing probability based on a previous spike at that lag, is shown in red whereas modulation lower than 1, representing a decrease in firing probability, is shown in blue. Visually, we find that the estimated trajectories of the modulation values are very close to the true simulated values at all lags and times, although the estimates show some noisy fluctuations, especially at the higher time lags. By considering only those points whose point-wise significances achieved a p-value below 0.05, the noisy changes at higher lags disappear as shown in the middle panel. The important dynamic features of the simulation model relating to the attenuation of the negative modulation around the 30 ms lag and of the positive modulation at the 50 ms lag are well captured.

To assess the confidence of the state-space adaptive model, we tracked the standard deviations of parameter estimates associated with the baseline intensity, bursting propensity (2 ms lag), and beta rhythmic firing (30 and 50 ms lags) through time. Figure 2.5C displays the estimated standard deviations of these parameters. They all show relatively small and stable standard deviations, indicating a high degree of confidence in the estimates through time. The estimated standard deviation is large at the beginning and the end of the trial, because the smoothing algorithm has less data with which to compute the estimates at these points. However, within 100 ms the variability quickly drops to a consistent value. Additionally, we find that the estimated confidence in the parameters related to the 30 and 50 ms lags is much smaller than those related to the baseline or bursting parameters, suggesting that these dynamics can be tracked with a high degree of confidence.

Figure 2.5D shows a K-S plot of the rescaled ISIs for the estimated state-space adaptive model. The K-S plot follows the 45 degree line closely, with the majority of the larger rescaled ISIs falling inside the 95% confidence bounds. Similar to the K-S

plot for the two-state model, the largest deviations from the smallest ISIs perhaps indicates some minor model misfit associated with the refractoriness and bursting of the STN neurons, which disappear if simulated at a higher temporal rate (not shown).

2.3.5 Tracking changes in spiking dynamics—STN data

Next, we applied the adaptive model to the spike train data recorded in STN from Parkinson’s patients. Figure 2·6 shows the results for a representative, rhythmically firing neuron. Figure 2·6A shows the adaptive estimate of the baseline firing intensity, along with pointwise 95% confidence bounds, at each time. Visually, we can identify an increase in the baseline firing intensity that begins around 400 ms before the initiation of the hand movement. The baseline intensity keeps rising up until the movement begins, remains relatively stable for around 300 ms, and then decreases back to the pre-movement level by around 800 ms after movement began.

Figure 2·6B shows the estimated history dependent modulation of firing intensity as a function of time relative to movement onset on the x-axis and lag between 10-100 ms on the y-axis. The top panel shows the estimated modulation parameter values at each lag and time point. The bottom panel shows only those parameters whose point-wise significance, computed using the estimated posterior covariance matrix, attained a p-value less than $p = 0.05$, with the rest of the modulation values set to 1. The large blue regions between lags of 20 ms to 35 ms at the beginning and end of the trial represent negative modulation, or a decrease in the firing probability as a function of previous spiking at these lags, which is attenuated starting around 1100 ms before the onset of hand movement and returns around 800 ms after movement onset. Notice that this change is detectable well in advance of the change in the baseline firing rate, which does not begin until around 400 ms prior to movement onset. The large red regions between lags of 40 ms and 60 ms indicate positive modulation, or an increase in instantaneous firing probability as a function of previous spikes at these

lags, which is attenuated around 400 ms before movement onset and returns to its initial level around 800 ms after movement onset. The time course of the change in this parameter is more similar to that of the baseline rate than the 30 ms modulation parameter.

This suggests a dynamic pattern of transition between spiking states, whereby the change in rhythmic spiking at beta frequencies initially comes about by a reduction in inhibitory control of spiking, which is followed by reduction in excitatory control at the beta period and a simultaneous increase in baseline firing rate. The top panel also shows short noisy changes in the estimated modulation parameters at higher time lags, which disappear when considering modulation estimates below the point-wise 5% significance level, as shown in the bottom panel, consistent with the simulation study.

Figure 2·6C shows the estimated standard deviations for the background rate and modulation parameters at lags of 2 ms, 30 ms and 50 ms. The standard deviation of the baseline firing rate shows dynamic behavior that mimics that of the estimated mean baseline rate. This is not surprising; the estimated mean and variance of the baseline rate of a point process are typically associated. The standard deviation for the modulation parameter at a 2 ms lag related to bursting is small relative to the mean modulation that takes values near 1. The standard deviation is minimal during the movement, where bursting behavior decreases. As with the simulation study, the standard deviation of the parameters related to the rhythmic firing at beta frequencies is small compared to other lags, suggesting that these dynamics can be tracked with a high degree of confidence.

Figure 2·6D shows the K-S plot of the time-rescaled ISIs for the adaptive model. The model fits well and provides a slight improvement over the two-state model shown in Figure 2·3. There is still a small but significant lack of fit associated with

the smallest rescaled ISIs. This suggests that there may still be some unmodelled features of the dynamics not represented in these models. For example, the model we used does not account for changes in firing as a function of movement direction. Typically, model identification for point process spiking models involves an iterative search across multiple classes of models incorporating different covariates and functional forms. Here, we have presented a simplified model to highlight changes in the rhythmic spiking dynamics rather than attempting to find a single model that completely describes all features of the spiking structure in the data.

2.4 Discussion

Increasingly, researchers studying the electrophysiological origins and mechanisms of various neurological diseases are acknowledging the importance of characterizing the dynamic properties of neural rhythms. The most prevalent methods for studying such dynamics typically focus on continuous brain signals such as LFPs, EEG, MEG, and the like. When applied to spike train data, these methods can lead to results that are difficult to interpret, or inappropriate because of incorrect modeling assumptions. Here we have presented a state-space point process framework that allows us to first identify dynamics associated with rhythmic spiking and then characterize the transitions between distinct dynamic states.

Specifically, we applied the point process modeling framework to characterize the transition of history-dependent firing activity in STN during a hand movement task. Model estimates showed that the effect of past spiking on the instantaneous probability of firing a spike were small on a single trial, but easily detectable across trials for dataset of the size typically collected in these experiments. We showed that a maximum likelihood ratio test was able to identify very subtle but statistically significant changes in these dynamics.

Furthermore, using a state-space adaptive algorithm, we were able to track and characterize the time course of the changes in modulation of the firing intensity as a function of past spiking history. Although the trajectory of the parameter estimates varied slightly depending on the regions of STN where the signals were recorded, we observed consistent trends across neurons in the time course of the modulation. The modulation of the baseline firing rate on average began to increase around 500 ms before movement onset, reached a maximum around 150 ms before movement onset, fluctuated around the peak value until 100 ms after movement onset, and finally returned to its original value until about 700 ms after movement onset. The parameter related to modulation due to past spiking 50 ms in the past followed a similar time course. However, the estimated influence of past spiking at a 30 ms lag showed more consistent dynamics across neurons and had a temporal trajectory that began to attenuate at earlier times in the planning period, suggesting that movement initiation begins with a release of inhibitory control. These results corroborate previous physiological findings that have characterized the role of the basal ganglia in the disinhibition of selected motor programs (Mink and Thach, 1993; Mink, 1996).

The proposed framework has several important features. First, by proposing an explicit point process model we obtain estimators that have clear interpretations in terms of the probability of observing a spike at any time. By working in the time domain, we need not assume a stationary spiking process, as is common in frequency domain analyses. We also avoid common errors in interpretation that can occur when frequency domain analyses fail to consider the effects of point process data on power spectral estimates across all frequencies (Brillinger, 1975).

Second, the point process modeling framework provides a great deal of flexibility in assessing relationships between the instantaneous probability of spiking, the past spiking history, and other covariates that may influence neural spiking. In this case,

we chose to model the effect of past spiking history using cubic spline basis functions. This model formulation allowed us to examine the effects of multiple temporal lags simultaneously with a relatively small number of degrees-of-freedom. This has the advantage of reducing the chance of overfitting in our models, improving the computational efficiency of the estimation algorithms, and increasing the statistical power of the hypothesis tests.

Third, the point process modeling framework allows us to examine the effects of multiple model components simultaneously. Simple visualization methods such as autocorrelation function typically examine the effect of previous spiking at a single time lag and can therefore fail to account for influences of one covariate on the estimated influence of another (Bar-Gad et al., 2001). More recently, point process models have been used to study multiple features of rhythmic neural spiking in the STN of Parkinsons patients (Sarma et al., 2009; Sarma et al., 2010). These models similarly incorporate history dependent modulation to capture features such as bursting and beta band oscillations, and include task dependent components such as movement direction and information related to movement cues. However, these models assume constant parameter values and do not attempt to capture rhythmic neural dynamics as a function of movement planning and execution. Forth, the parametric modeling approach provided for a powerful hypothesis test for identifying subtle changes in the firing dynamics. The parametric approach in conjunction with the time rescaling theorem also provides a simple method for rapid simulation of data. We took advantage of this to compute the power of the maximum likelihood ratio hypothesis test. This will also be critical in expanding these analysis methods to study interactions between large neural ensembles (Brown et al., 2002).

Finally, the state-space smoothing algorithm allowed us to track the time course of the change in history dependence, and therefore characterize a critical transition

in rhythmic spiking dynamics. We were able to maintain adaptive estimates of both the parameter mean and variance-covariance structure, permitting the construction of confidence intervals and the identification of interactions between covariates (Eden et al., 2004).

We foresee several potential improvements and extensions of this framework. The algorithm we described can be applied to a single trial of data or to multiple trials that share common dynamical properties. Here we assumed that the spiking dynamics were identical across all trials. This allows us to combine data over trials to detect very subtle changes in the modulation structure through time. If the changes to the modulation parameters as a function of movement are large, then we expect to be able to detect those changes using a single trial of data. For the STN data we analyzed in this paper, we found that movement is associated with very subtle changes in the modulation parameters. In this case, assuming these changes are consistent across trials allows us to detect these small changes with confidence. If the assumption that these changes are consistent across trials is not true, then our model is misspecified, which will affect the bias and variance characteristics of the estimates. In general, we would expect such misfit to manifest as a loss of statistical power, making the identification of significant changes more difficult, and to manifest as a reduction in the goodness-of-fit measures. While no model is completely correct, the fact that our models capture the structure of the data well suggests that the effect of the model misspecification on the analysis results is not substantial.

Here we used a simplified model of neural spiking to illustrate the problem of identifying and tracking transitions in rhythmic spiking dynamics. In a more complete study, we would develop more accurate and physiologically realistic models by iteratively refining the model class to incorporate other covariates or other functional relationships between covariates. The question of relating statistical models that de-

scribe the structure of neural spiking patterns and physiologically realistic models that describe the mechanisms of spike generation is one of increasing importance. As our models become more complex, estimation and testing methods will need to adapt to accurately identify dynamic transitions. Finally, here we limited ourselves to modeling the spiking activity of a single neuron because in our experimental dataset very few neurons were simultaneously recorded. Typically, neural ensembles work together to generate and maintain rhythmic activity. We can extend the point process framework by modeling the firing probability of each neuron based on its own history as well as the activity of other simultaneously active neurons in the same or in different brain regions. The methods we developed here can be easily extended to address issues of synchronization by incorporating the past history of simultaneously recorded neurons with parameters related to coupling between neurons into the conditional intensity model with parameters related to coupling between neurons.

As we learn more about the electrophysiological processes underlying neurological disease, the role of rhythms and transitions will continue to grow. One important front in our pursuit to understand these processes is the development of useful statistical estimation and inference procedures—procedures that respect the structure of the data, that rely on appropriate modeling assumptions, and that provide a high degree of statistical power. Point process methods, such as the ones explored here, provide a powerful approach for understanding dynamic structure in spike train data, both in the healthy and diseased brain.

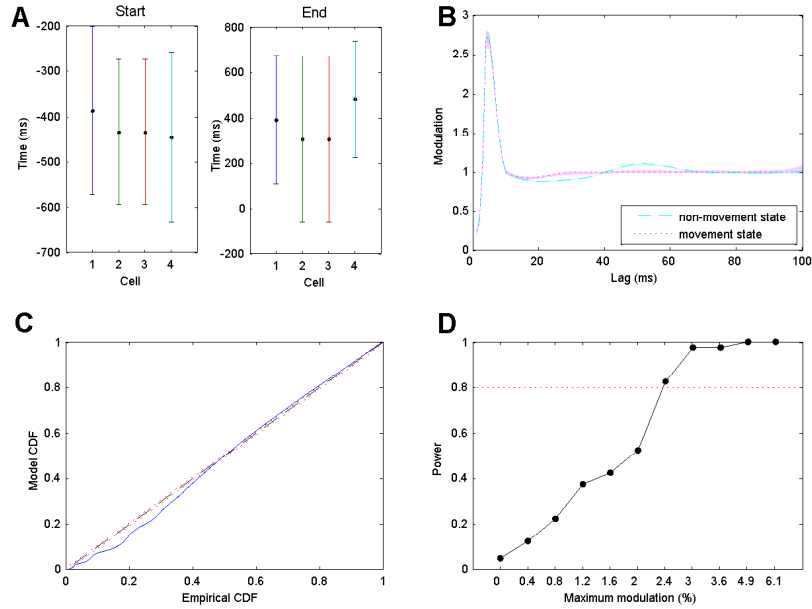


Figure 2.2: (A) Estimates and 95% confidence bounds for estimated transition times for the start and end points of the movement spiking state for simulated spike trains of four neurons under the two-state model. Actual transitions occurred at -400 ms and 400 ms for start and end transitions, respectively. (B) 95% confidence bounds for the estimated history dependent modulation parameters as a function of lag and the true parameter values for simulated spike data from a single neuron. Cyan dashed line and magenta dotted line represent true modulation during non-movement state and during movement state, respectively. The transparent regions in cyan and magenta correspond to the estimated 95% confidence regions. These regions are narrow and tend to contain the true value at every lag. (C) K-S plot comparing empirical and model CDFs of rescaled simulated ISIs from estimated point process model. (D) Power of the maximum likelihood ratio test as a function of the modulation of the history parameters at 30 ms and 50 ms lags between the non-movement and movement states. The test detects as little as 2.4% modulation in these parameters with 82.5% probability.

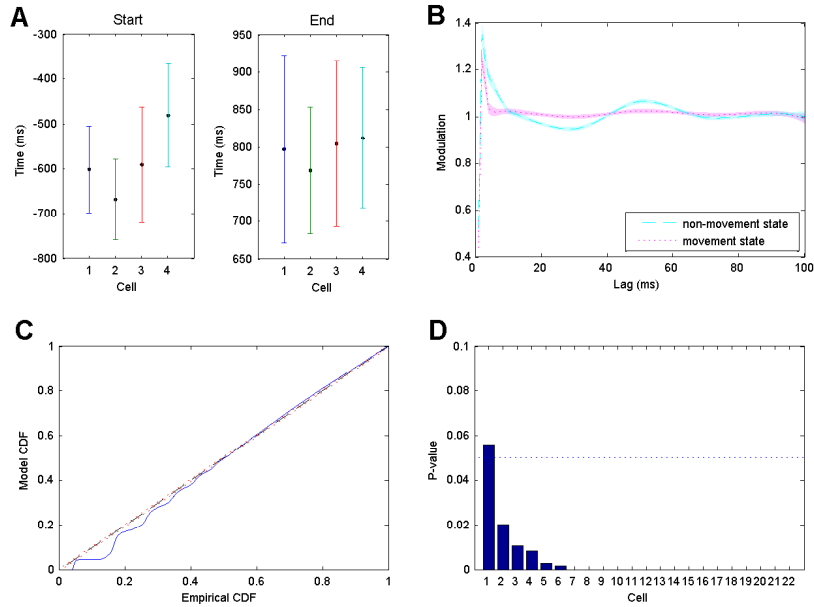


Figure 2.3: (A) Estimates and 95% confidence bounds for transition points between non-movement and movement states for four representative neurons. (B) Estimated history dependent modulation parameters as a function of lag with 95% confidence bounds for a single representative neuron. Cyan dashed line and magenta dotted line represent modulation estimates during non-movement state and during movement state, respectively. Transparent regions represent corresponding the 95% confidence regions. (C) K-S plot comparing empirical and model CDFs of rescaled ISIs from estimated point process model. (D) P-values for the maximum likelihood ratio test for significant changes in modulation between states for all 22 cells, in descending order.

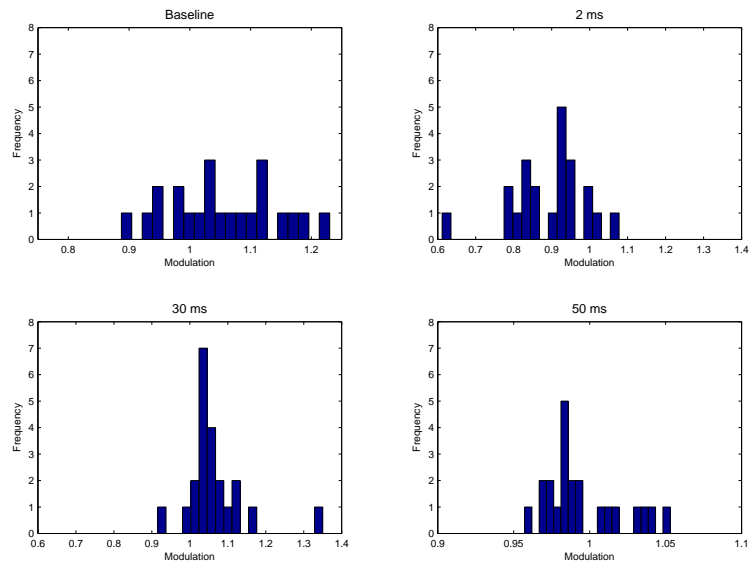


Figure 2.4: Histograms of ratio of estimated model parameters for baseline intensity and modulation at lags of 2, 30, and 50 ms across all 22 neurons analyzed. A value of 1 indicates no change in modulation between the two states, a value above 1 indicates an increased estimate in the movement state, relative to the non-movement state, and a value below 1 indicates a decreased estimate in the movement state.

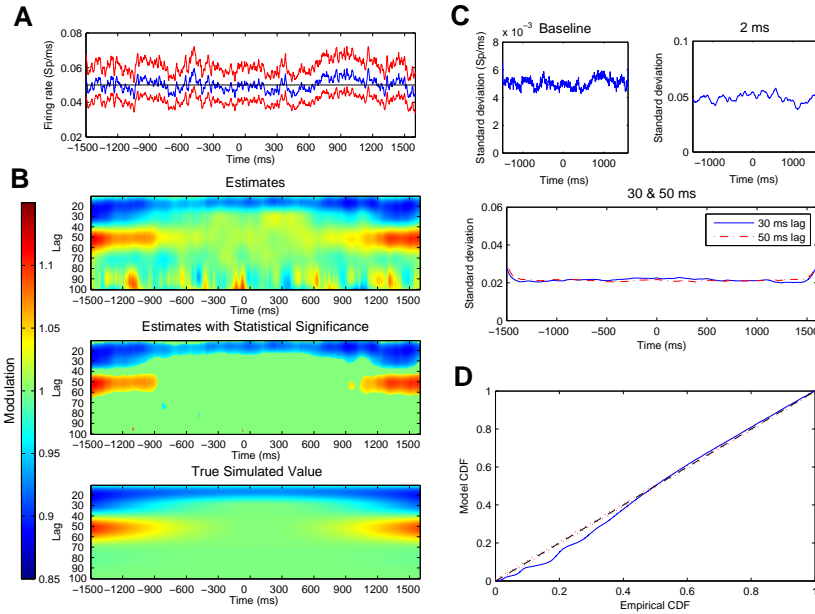


Figure 2-5: (A) Estimated trajectory and uncertainty of baseline firing rate parameter for simulated spiking data. Blue line represents the estimates and red lines denote 95% confidence regions at each time. Black line represents the true constant baseline firing rate parameter used to simulate the data. (B) Estimated and true values of history dependent modulation of firing intensity as a function of time relative to movement onset on the x-axis and lag between 10–100 ms on the y-axis. (C) Standard deviations for parameter estimates for baseline firing rate, and modulation at lag 2 ms, 30 ms and 50 ms lags. In each case, estimated standard deviations are small compared to estimated mean parameter values, indicating high confidence in estimates. (D) K-S plot comparing empirical and model CDFs of rescaled simulated ISIs from estimated point process model.

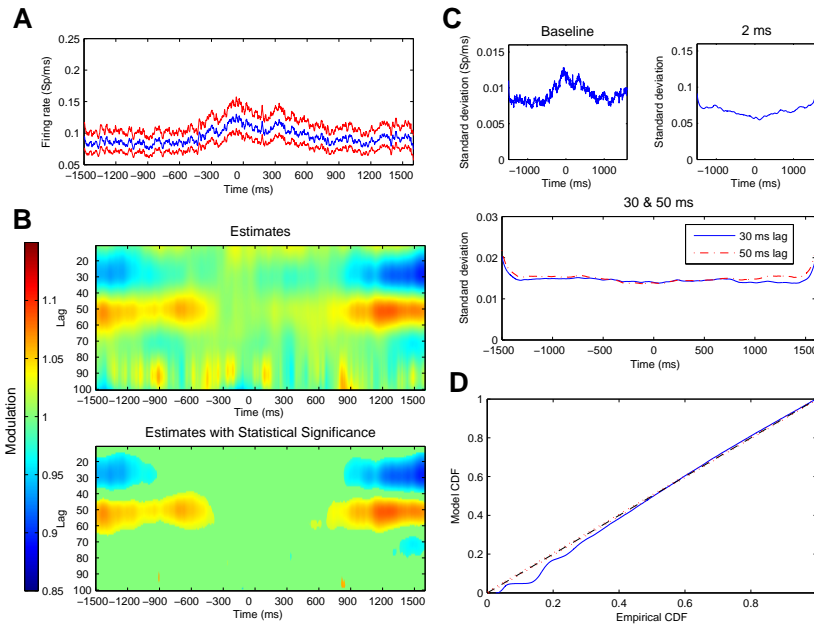


Figure 2-6: (A) Estimated trajectory and uncertainty of baseline firing rate parameter for observed spiking data from a representative STN neuron. Blue line represents the estimates and red lines denote 95% confidence regions at each time. (B) Estimated values of history dependent modulation of firing intensity as a function of time relative to movement onset on the x-axis and lag between 10–100 ms on the y-axis. Blue regions at beginning and end of trial from 20 to 35 ms lags indicate regions of negative modulation and red regions at beginning and end of trial from 40 to 60 ms lags indicate regions positive modulation. (C) Standard deviations for parameter estimates for baseline firing rate, and modulation at lag 2 ms, 30 ms and 50 ms. (D) K-S plot comparing empirical and model CDFs of rescaled ISIs from estimated point process model.

Chapter 3

Clusterless Decoding of Position From Multiunit Activity Using A Marked Point Process Filter

3.1 Introduction

Neural systems encode information about external stimuli in temporal sequences of action potentials. Because action potentials have stereotyped, impulse waveforms, they are most appropriately modeled as point processes (Brillinger, 1992). Neural systems are moreover dynamic in that the ensemble firing of populations of neurons, representing some biologically relevant variable, continually evolves. Decoding algorithms based on adaptive filters have been developed to study how the firing patterns maintain dynamic representations of relevant stimuli. More specifically, both discrete-time and continuous-time point process filter algorithms have been applied with great success to address problems of estimating a continuous state variable (Eden et al., 2004; Smith and Brown, 2003; Smith et al., 2004), such as the location of an animal exploring its environment (Brown et al., 1998; Huang et al., 2009; Koyama et al., 2010).

A prerequisite for these increasingly efficient decoding methods is the application of spike-sorting: the waveforms recorded extracellularly at electrodes must be clustered into putative single neurons. Therefore, the accuracy of the spike-sorting critically impacts the accuracy of the decoding (Brown et al., 2004). Many algorithms

for spike-sorting, whether real-time and automatic or offline and manual, have been developed since the 1960s (Lewicki, 1998; Wild et al., 2012). The majority of these algorithms are clustering-based methods, allocating each spike to a putative single cell based on the characteristics of spike waveforms. These types of pure waveform, hard boundary spike-sorting algorithms, suffer from many sources of error such as nonstationary noises, non-Gaussian clusters, and synchrony (Lewicki, 1998; Harris et al., 2000; Quiroga, 2012). In addition, they also have been shown to yield biased and inconsistent estimates for neural receptive fields (Ventura, 2009b). Another clustering method, soft or probabilistic spike assignment, has been incorporated into some decoding paradigms and analyses have shown that these algorithms can yield unbiased estimates of stimulation parameters (Ventura, 2008; Ventura, 2009a). Nonetheless, these algorithms, like most hard sorting methods, are not well suited to real-time implementations, both because the algorithms are too computationally intensive and because they rely on having the entire dataset available for the clustering algorithm.

More recently, decoding methods without a spike-sorting step have been investigated (Luczak and Narayanan, 2005; Stark and Abeles, 2007; Fraser et al., 2009; Chen et al., 2012; Kloosterman et al., 2014b). Chen et al. (2012) and Kloosterman et al. (2014) developed a spike feature decoding paradigm for unsorted spikes using a time-homogeneous spatio-temporal Poisson process. It incorporates a covariate-dependent method to estimate a nonparametric distribution of the animal's position. This improves decoding performance by using information that is otherwise excluded by sorting spikes into discrete clusters. However, this method does not incorporate a model of the animal's position trajectory, and therefore the decoding results can depend substantially on specific model parameters such as the choice of a discrete time bin width: if the time scale is too broad, the algorithm cannot track the stimuli fast enough; if too narrow, it cannot integrate information provided by spikes nearby

in time. Additionally, this algorithm extracts information from spike intervals, but does not optimally incorporate information from intervals that contain no spikes. In the Extensions and Discussion to Chen et al. (2012), they suggested the possibility of applying a state-space framework with a temporal prior, but do not provide a complete algorithm or an implementation of this method.

To address these issues, we generalize and extend this decoding paradigm for unsorted spikes to allow for general marked point process models. In doing so, we develop an iterative algorithm that solves the marked point process filter problem. In particular, this allows us to model neural activity that is dependent on the timing and mark values of previous spikes. In the current work, we propose a novel algorithm for adaptive decoding of spiking activity which avoids the clustering problem of spike-sorting entirely by defining a joint model for the spike waveform and receptive field structure, and uses a state-space model to incorporate knowledge of the properties of the signal to decode. The resulting algorithm is a general marked point process filter.

The goal of the proposed algorithm is to decode an adaptive state variable, in this case the dynamic trajectory of the rat along a track, by computing the posterior distribution of the state process conditioned on the set of observations up until the current time. Our algorithm takes directly as inputs the recorded spikes where each spike is associated with a vector of characteristic features selected from the raw waveforms. Such inputs can be mathematically described by a marked point process, where the points are the spike times and the marks are their corresponding waveform. In their seminal book on point processes in 1980, D. R. Cox and V. Isham have already explicitly suggested that “in the neurophysiological example mentioned [...] the mark could be the magnitude of the peak signal at the point in question” (Cox and Isham, 1980).

In this paper, we first characterize the spiking activity of an ensemble of neurons

using the conditional intensity function for marked point processes. Next, we construct a stochastic state marked point process filter to iteratively calculate the full posterior probability for the state variable. We illustrate our approach in a simulation study where the decoding algorithm is used to reconstruct an animal’s position from unsorted multiunit place cell spiking activity. We also apply the algorithm to experimental data recorded in the hippocampus of a rat navigating a linear environment. We then compare the quality of fit of our clusterless decoding algorithm to that of a traditional spike-sorting and decoding algorithm.

3.2 Methods

Any orderly point process can be fully characterized by its conditional intensity function (Daley and Vere-Jones, 2003). A conditional intensity function describes the instantaneous probability of observing a spike, given previous spiking history. By relating the conditional intensity function to specific biological and behavioral signals we can specify a spike train encoding model. The conditional intensity function also generalizes to the marked case, in which a random vector, termed a mark, is attached to each point. In the case of tetrode recordings, for example, the mark could be a length four vector of the maximum amplitudes on each of the four electrodes at every spike time. A marked point process is completely defined by its joint mark intensity function

$$\lambda(t, \vec{m}|H_t) = \lim_{\Delta \rightarrow 0} \frac{\Pr(\text{a spike with mark vector } \vec{m} \text{ in } (t, t + \Delta]|H_t)}{\Delta}, \quad (3.1)$$

where H_t is the history of the spiking activity up to time t . $\lambda(t, \vec{m}|H_t)$ represents a joint stochastic model for the marks as well as the arrivals of the point process.

The joint mark intensity function characterizes the instantaneous probability of observing a spike with mark \vec{m} at time t as a function of factors that may influence

spiking activity. We posit that the spiking activity depends on some underlying internal state variable $x(t)$, such as an animal's location in space, that varies across time. We can therefore model the spiking activity as $\lambda(t, \vec{m}|H_t) = g(x(t), \vec{m}|H_t)$. When taking an integral of $\lambda(t, \vec{m}|H_t)$ over the mark space \mathcal{M} , we get the probability of observing a spike regardless of the mark value:

$$\Lambda(t|H_t) = \int_{\mathcal{M}} \lambda(t, \vec{m}|H_t) d\vec{m}. \quad (3.2)$$

$\Lambda(t|H_t)$ can be understood as the conventional conditional intensity function of a temporal point process. For a marked point process, $\Lambda(t|H_t)$ is called the intensity function of the ground process (Daley and Vere-Jones, 2003). The mark space \mathcal{M} can be of any dimension.

The goal of our decoding algorithm is to compute, at each time point, the posterior distribution of the state variable given observed marked spiking activity. To do this, we apply the theory of state-space adaptive filters (Haykin, 1996). Recursive filter equations can be solved in both discrete-time and continuous-time (Eden et al., 2004; Eden and Brown, 2008). In this paper we present an algorithm in discrete-time.

To describe the discrete-time filters, we partition an observation interval $[0, T]$ into $\{t_k : 0 \leq t_1 < \dots < t_N \leq T\}$ and let $\Delta_k = t_k - t_{k-1}$. The posterior density for the state variable at time t_k can be derived simply using Bayes' rule,

$$p(x_k | \Delta N_k, \vec{m}_k, H_k) = \frac{p(x_k | H_k) p(\Delta N_k, \vec{m}_k | x_k, H_k)}{p(\Delta N_k, \vec{m}_k | H_k)}, \quad (3.3)$$

where $x_k = x(t_k)$ is the state variable at time t_k , ΔN_k is the number of spikes observed in the interval $(t_{k-1}, t_k]$, and H_k is the spiking history up to time t_k . \vec{m}_k represents a collection of mark vectors \vec{m}_{k_i} , $i = 1, \dots, \Delta N_k$, observed in the interval $(t_{k-1}, t_k]$.

The first term in the numerator on the right hand side of equation (3.3), $p(x_k | H_k)$,

is the one-step prediction density defined by the Chapman-Kolmogorov equation as

$$p(x_k|H_k) = \int p(x_k|x_{k-1})p(x_{k-1}|\Delta N_{k-1}, \vec{m}_{k-1}, H_{k-1}) dx_{k-1}. \quad (3.4)$$

Here we have assumed that given the past state value, x_{k-1} , the distribution of the current state does not depend on past spiking activity. The integral in equation (3.4) typically cannot be solved analytically, but multiple numerical and approximation methods are available to compute its value at each time point. In this case, we performed numerical integration over a one-dimensional state space using a simple Riemann sum. If the state variable is high-dimensional, alternative methods can be considered to improve computational efficiency. In particular, when prior knowledge allows us to assume unimodality of the posterior density, a linear recursive Gaussian approximation to the posterior density can be constructed (Brown et al., 1998; Eden et al., 2004; Smith and Brown, 2003). Alternatively, when the posterior distribution is unknown, sequential Monte Carlo methods, also called particle filters, can provide efficient estimation (Doucet et al., 2001; Ergun et al., 2007). We show the construction of these two types of algorithms in the Appendix.

Equation (3.4) has two components: $p(x_k|x_{k-1})$, which is given by a state evolution model under the Markovian assumption, and $p(x_{k-1}|\Delta N_{k-1}, \vec{m}_{k-1}, H_{k-1})$, which is the posterior density from the last iteration step. We multiply this probability density function with the posterior distribution of the state variable at the previous time step t_{k-1} , and numerically integrate the product over all possible values of the previous state, x_{k-1} . The resulting integral is the a one-step prediction density at the current time t_k .

The second term in equation (3.3), $p(\Delta N_k, \vec{m}_k|x_k, H_k)$, is the likelihood or observation distribution at the current time, and can be fully characterized by the joint

mark intensity function $\lambda(t, \vec{m}|H_t)$:

$$p(\Delta N_k, \vec{m}_k | x_k, H_k) \propto \begin{cases} \exp[-\Delta_k \Lambda(t_k | H_k)], & \Delta N_k = 0; \\ \prod_{i=1}^{\Delta N_k} [\lambda(t_k, \vec{m}_{k_i} | H_k) \Delta_k] \exp[-\Delta_k \Lambda(t_k | H_k)], & \Delta N_k > 0. \end{cases} \quad (3.5)$$

We can interpret equation (3.5) by separating the product on the right hand side into two terms. The $\prod_{i=1}^{\Delta N_k} [\lambda(t_k, \vec{m}_{k_i} | H_k) \Delta_k]$ term characterizes the distribution of firing ΔN_k spikes, such that the mark value of the i^{th} spike in the interval $(t_{k-1}, t_k]$ is m_{k_i} , where $i = 1, \dots, \Delta N_k$. If no spike occurs, i.e., $\Delta N_k = 0$, this term equals 1. The $\exp[-\Delta_k \Lambda(t_k | H_k)]$ term gives the probability of not firing any other spikes in the observation interval, where $\Lambda(t_k | H_k) = \int_{\mathcal{M}} \lambda(t_k, \vec{m}_k | H_k) d\vec{m}_k$ as defined in equation (3.2).

The discrete-time likelihood or observation distribution defined by equation (3.5) assumes that within a small time step Δ_k , conditional on the history H_k and the current value of the state vector x_k , spiking activity for the neural ensemble is independent. However, the spiking activity at a given time step can explicitly depend on the past history of activity from the entire population, including dependencies between neurons.

The observation distribution is then multiplied by the one-step prediction density to get the posterior density at the current time. Note that we can drop the normalization term, $p(\Delta N_k, \vec{m}_k | H_k)$ because it is not a function of x_k . Substituting equation (3.4) into equation (3.3) yields a recursive expression for the evolution of the unnormalized posterior density:

$$p(x_k | \Delta N_k, \vec{m}_k, H_k) \propto p(\Delta N_k, \vec{m}_k | x_k, H_k) \int p(x_k | x_{k-1}) p(x_{k-1} | \Delta N_{k-1}, \vec{m}_{k-1}, H_{k-1}) dx_{k-1}. \quad (3.6)$$

3.3 Simulation Study

We first tested this approach on simulated data corresponding to the activity of two “place cells” firing as an animal traverses a linear track. Place cells are neurons in the hippocampus that are active primarily when an animal is located in a certain portion of its environment (O’Keefe and Dostrovsky, 1971; O’Keefe, 1979). Substantial amount of information about the position carried by place cells have been reliably quantified with a formal statistical algorithm and used to reconstruct the trajectory and predict the future position of the rat (Muller and Kubie, 1989; Wilson and McNaughton, 1993; Zhang et al., 1998; Brown et al., 1998).

In this section, we apply the clusterless decoding algorithm to the problem of decoding a one-dimensional position of a rat on a linear track using a marked point process arising from two hippocampal place cells, where the one-dimensional mark represents the peak height of the spike waveform. Here we note that accurate decoding of position based on place cell activity normally requires many cells. Our goal here is therefore not to test the accuracy of decoding, but rather to provide intuition about how our approach works in a simple case. First, we simulate the rat’s trajectory using an autoregressive process and the joint mark intensity function of the two cells using a mixture of two bivariate Gaussian distributions. Then, we use marked point process filter to reconstruct the location of the rat at each time step. We show that the filter yields exact full posterior densities that are often multimodal. We also reconstruct the trajectory using a traditional algorithm in which decoding was applied after spike-sorting. Last, we compare the performance of the clusterless decoding algorithm to the traditional decoding algorithm.

3.3.1 Data simulation

We simulated the movement of a rat running back and forth along a linear track with the following transition probability

$$p(x_k|x_{k-1}) \sim N(\alpha x_{k-1}, \sigma), \quad (3.7)$$

where we set $\alpha = 0.98$ and $\sigma = 0.05$ to generate an autoregressive process whose steady-state standard deviation is 1.26.

We defined the joint mark intensity of two hippocampal place cells using a bivariate Gaussian mixture function,

$$\lambda(t, \vec{m}) = \sum_{c=1}^2 \exp \left\{ \alpha^c - \begin{pmatrix} x(t) - \mu_x^c \\ \vec{m} - \mu_{\vec{m}}^c \end{pmatrix}^T \begin{bmatrix} 2(\sigma_x^c)^2 & 0 \\ 0 & 2(\sigma_{\vec{m}}^c)^2 \end{bmatrix}^{-1} \begin{pmatrix} x(t) - \mu_x^c \\ \vec{m} - \mu_{\vec{m}}^c \end{pmatrix} \right\}. \quad (3.8)$$

This model assumes that spikes originate from an ensemble of two cells and that the joint mark intensity function for each cell's place field and marks is a multivariate Gaussian function. In other words, each cell has a place field with a Gaussian shape, and the marginal distribution of the marks given the rat's position is a multivariate Gaussian distribution.

The parameters for this model include α^c for $c = 1, 2$, which controls the maximum in-field firing rate for the c^{th} neuron. $x(t)$ is the animal's linearized position at time t and \vec{m} is the one-dimensional mark value. μ_x^c is the center of the place field and $\mu_{\vec{m}^c}$ is the mean of the density function for marks for neuron c . σ_x^c and $\sigma_{\vec{m}}^c$ are the standard deviations of the place field and the mark space for neuron c respectively. In this illustrative example, we set the centers of the place field at $\mu_x^1 = -1.5$ and $\mu_x^2 = 1.5$ respectively, with a $(\sigma_x^1)^2 = (\sigma_x^2)^2 = 0.1$ variance. We set the maximum in-field firing rate of both neurons to $\alpha^1 = \alpha^2 = 100$ spikes per second. The one-dimensional mark

space for the two neurons are centered at $\mu_{\vec{m}}^1 = 10$ and $\mu_{\vec{m}}^2 = 13$ respectively, with standard deviations $\sigma_{\vec{m}} = \sigma_{\vec{m}}^1 = \sigma_{\vec{m}}^2$ varying between 0.01 and 5.

By plugging in the simulated position trajectories, $x(t)$, into Equation (3.8), we computed the instantaneous intensity at each time step. The underlying unmarked spike trains are generated under an inhomogeneous Poisson process model with rate $\Lambda(t|H_t)$ given by equation (3.2). The marks associated with the spike trains are sampled from the probability density $f(\vec{m}|t) = \frac{\lambda(t, \vec{m})}{\Lambda(t|H_t)}$.

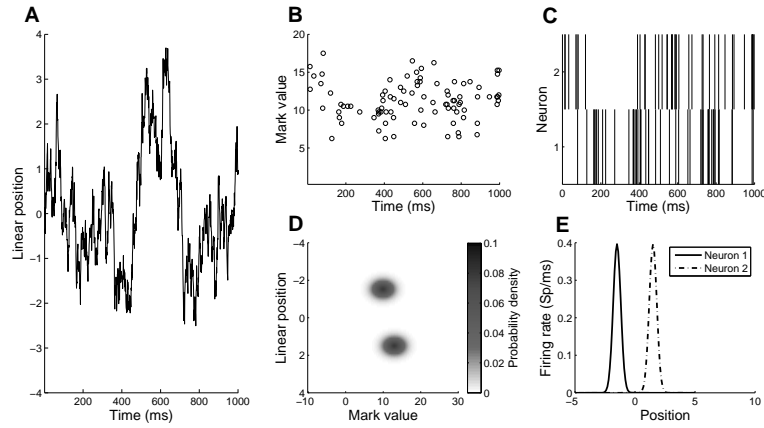


Figure 3-1: (A) Simulated trajectory of the animal running back and forth on a linear track. (B) One second unsorted spike train with marks from two simulated neurons. (C) One second raster plot of spikes without marks where the spikes were sorted via a linear discriminant function on their mark values. (D) True joint mark intensity of the two simulated neurons as a function of linear position on the y-axis and mark value on the x-axis. The place fields center at -1.5 and 1.5 respectively, with a 0.1 variance. The mark spaces are one-dimensional with mean of 10 and 13 respectively, and standard deviation of 2. (E) True rate of each of the two simulated neurons as a function of linear position. It represents the place field of each neuron.

Figure 3-1 displays the results of a representative, one-second trial where spiking activity of two neurons with moderate mark space overlap $\sigma_{\vec{m}} = 2$ was simulated. Figure 3-1A shows the simulated trajectory of the animal with the linear location on the y-axis. Figure 3-1B plots the simulated spike train with the mark value on the

y-axis. Visually, we cannot identify a clear clustering of the mark values. Figure 3.1D shows the true joint mark intensity function $\lambda(t, \vec{m})$ used to simulate the marked spike train in Figure 3.1B. The two simulated neurons overlap over a moderate amount of mark space, but are fairly localized in space, with little overlap of their place fields.

3.3.2 Decoding results

In this subsection, we present the decoding results of the simulated data described in subsection 4.3.1 Data Simulation using the clusterless decoding algorithm described in the Methods section. When calculating the observation distribution defined in Equation (3.5), we used the true joint mark intensity function specified in Equation (3.8). In real data, we will need to estimate an encoding model for the joint mark intensity function. However, this simulation study affords us the opportunity to explore errors due purely to the decoding algorithms.

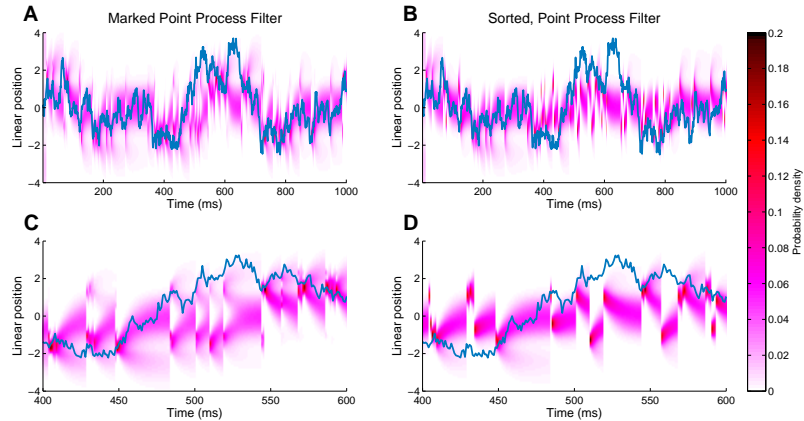


Figure 3.2: Decoding results for two simulated neurons. (A) Posterior density of the animal’s linear position using the clusterless decoding algorithm. The blue line represents the actual position of the animal. The red represents the posterior density at each time step. (B) Posterior density of the animal’s linear position using the decoding with spike-sorting. (C) Clusterless decoding results zoomed in at time between 400 and 600 ms, showing the multimodality of the posterior around 430 ms. (D) Decoding with spike-sorting results zoomed in at time between 400 and 600 ms.

Figure 3·2A shows the posterior density for the position of the animal during a one-second trial as a function of time, using the clusterless decoding algorithm. The blue line shows the true position of the animal, and the red region is the estimated posterior density at that time step. We can see that the estimate tracks the true position relatively closely, and the region of high posterior density covers the true position most of the time.

In order to compare with the traditional approach where decoding is applied after spike-sorting, we also implemented a decoding algorithm on individual neurons after classifying the spikes into clusters using linear discriminant functions (Bishop, 2006). Here we assume the number of neurons being recorded is known and that we have true knowledge of the rate function $\lambda^c(t)$ about neuron c . This is more knowledge than is typically available for clustering and decoding, where both the number of neurons and place fields must be estimated. The time-dependent rate function for individual neuron, $\lambda^c(t)$, takes the following form:

$$\lambda^c(t) = \exp \left\{ \alpha^c - \frac{1}{\sqrt{2(\sigma_x^c)^2}} (x(t) - \mu_x^c)^2 \right\}. \quad (3.9)$$

Note that $\lambda^c(t)$ represents the conventional rate function for an inhomogeneous Poisson process, not the joint mark intensity process we have used to define our marked point process model.

Figure 3·1C plots the spike raster of the two simulated neurons sorted using a linear discriminant function on their marks. Figure 3·1E shows the true condition intensity function $\lambda^c(t)$ for each neuron c , i.e., the place field of each neuron.

In this case, the unnormalized posterior density is of the form

$$p(x_k | \Delta N_k^1, \Delta N_k^2) \propto \prod_{c=1}^2 \Pr(\Delta N_k^c | H_k) \int p(x_k | x_{k-1}) p(x_{k-1} | \Delta N_{k-1}^1, \Delta N_{k-1}^2) dx_{k-1}. \quad (3.10)$$

The probability of seeing ΔN_k^c spikes from neuron c during the time interval $(t_{k-1}, t_k]$ is

$$\Pr(\Delta N_k^c | H_k) \propto (\lambda_k^c \Delta_k)^{\Delta N_k^c} \exp\{-\lambda_k^c \Delta_k\}, \quad (3.11)$$

where, for simplicity, λ_k^c denotes $\lambda^c(t_k)$.

Figure 3·2B and 4·2D are results using decoding with spike-sorting where the decoder is applied after classifying the spikes into two clusters. Figure 3·2B shows the posterior density using decoding with spike-sorting. Visually, we obtained comparable tracking for the animal's trajectory.

To illustrate the advantage of using a clusterless decoding algorithm when clear clustering of neurons is difficult, we zoomed in on the decoding results at times between 400 ms and 600 ms, shown in Figure 3·2C and 3·2D. From Figure 3·1B and 3·1C, we can see that even though the rat's position is positive, and therefore only Neuron 2 is spiking, the spike waveforms cannot be perfectly resolved. This leads to incorrect decoding results using the presorted spiking activity. However, the marked point process decoder results in a bimodal posterior distribution that accurately reflects the uncertainty due to the waveform overlap. Figure 3·2C demonstrates the multimodality of the posterior estimated by a clusterless decoder. For example, at time around 430 ms when a new spike arrives, the posterior density splits into two modes, one near a linearized position of -1 and another around a linearized position at 1.5. The posterior density is slightly higher in the region correctly predicting the animal's position in the negative regime. For the decoding with spike-sorting results

shown in Figure 3-2D, the algorithm has some trouble tracking the trajectory after new spikes from both neurons arrive because the posterior density in this case is not multimodal. For example at time around 430 ms, the posterior density incorrectly estimates the animal's position at a positive value near 1.

3.3.3 Goodness-of-fit analysis

In order to compare the quality of fit between the two decoding algorithms, we simulated 100 trials of spiking activity for different degrees of overlap between the two neurons in the mark space. We evaluate the ability of these two algorithms to track the location of the animal when the overlaps in the mark space between the joint mark intensity functions of the simulated neural ensemble vary from low ($\sigma_m^c = 0.01$) to high ($\sigma_m^c = 5$) in units of the number of standard deviations, using two different measures: the root-mean-squared error (rMSE) between the true positions of the animal and their estimated values and the percentage of the time that the true position values were covered by the 99% highest posterior density (HPD) region (Casella and Berger, 2001).

To calculate the 99% HPD region, we find the largest value y_0 such that

$$\int_{\{x:\rho_k(x)>y_0\}} \rho_k(x)dx = 0.99, \quad (3.12)$$

where ρ_k denotes the posterior density at time step t_k as defined in equation (3.3) for marked point processes and in equation (4.10) for the unmarked case. The 99% HPD region is thus given by $\{x : \rho_k(x) > y_0\}$. Because an HPD region indicates the region of highest posterior probability regardless of contiguity, it is a useful measure of uncertainty when the posterior density can be multimodal. If the state space model is correct, the 99% HPD region should contain the true position of the animal at any time step with probability 0.99.




Figure 3.3: Quality of fit comparison between the two decoding algorithms. (A) Root-mean-squared error (rMSE) between the true positions of the animal and their estimated values averaged across 100 trials as a function of the overlap between mark spaces of two simulated neurons. Error bars represent 2 standard deviations from the mean rMSE. The dotted line represents the performance of clusterless decoding method. The solid line represents the performance of decoding with spike-sorting method. (B) Fraction of time that the true position values were covered by the 99% highest posterior density (HPD) region averaged across 100 trials as a function of the overlap between mark spaces of two simulated neurons. Error bars represent 2 standard deviations from the mean coverage probability. The dotted line represents the performance of clusterless decoding method. The solid line represents the performance of decoding with spike-sorting method.

Figure 3.3A shows the rMSEs between the true positions of the animal and their estimated values averaged across 100 trials as a function of σ_m . Error bars represent 2 standard deviations from the mean rMSE. The dotted line represents the performance of the clusterless decoding method. The solid line represents the performance of the decoding with spike-sorting method. When the standard deviations of both mark spaces exceed 1, the clusterless decoder consistently gives the lower mean rMSE and there is no overlap between the error bars for the two algorithms.

The comparison of the rMSEs shows that as the overlap between the marks for the individual neurons increases and clustering of the spike waveforms becomes more

difficult, the clusterless decoding algorithm provides a consistently more accurate estimate of the rat's position than the decoding algorithm using presorted spikes.

Figure 3-3B shows the fraction of time that the true position values were covered by the 99% HPD region averaged across 100 trials as a function of σ_m , which measures how well each algorithm characterizes their uncertainty. The flat dotted line with narrow error bars illustrates that regardless of the degree of the overlap between mark spaces, the true trajectory stays in the 99% HPD region of the clusterless decoding method around 99% of the time with very little variance about this percentage of time across the 100 repeated trials of the simulation. The descending solid line with widening error bars illustrates that, for decoding with spike-sorting, as the degree of the overlap between mark spaces increases, the fraction of time that the true trajectory stays in the 99% HPD region is decreasing from 99% to 80% with an increasing variance.

The comparison of the 99% HPD region illustrates two advantages of the clusterless decoding algorithm regarding the uncertainty of the state estimates. First, the width of a 99% HPD describes the degree of uncertainty in the estimates. As the overlap between the mark spaces increases, the width of the 99% HPD region of the decoding algorithm using presorted spikes increases, indicating decreasing certainty in the estimated position. However, the clusterless decoding algorithm provides a narrower 99% HPD region regardless of the mark space overlap. This means that the clusterless decoding algorithm provides more estimates of the rat's position and is less influenced by the degree of overlap in the mark space. Second, the percentage of time within the 99% HPD represents the accuracy of uncertainty in the estimates. As the overlap in the mark distribution between the neurons increases, the decoding with presorted spikes estimates the uncertainty in position with less accuracy, while the estimated certainty of the clusterless decoding algorithm remains the correct at

around 99%. Therefore the estimated certainty for the clusterless decoding algorithm is both higher and more accurate than the estimated certainty for decoding with presorted spikes.

3.4 An Application to Position Decoding from Multiunit Activity in Rat Hippocampus

We also applied the clusterless decoding algorithm to experimental data recorded using a multi-electrode array in the hippocampus of a rat running back and forth on a linear track. The data used in this analysis were recorded from 5 tetrodes in the dorsal CA1 and CA2 regions of the hippocampus. Manual clustering of spike waveforms above a $40 \mu\text{V}$ threshold during the linear track session yielded fifteen distinct unit clusters. Four units with less than 100 spikes during the 840 second recording session were excluded. One of the remaining units was a putative interneuron, identified by firing rate exceeding 7 Hz and exceptionally narrow spikewidth. In the clusterless decoding, all of the thresholded spikes above the $40 \mu\text{V}$ threshold, including this putative interneuron, were included. Details of the experimental preparation, data acquisition and choice of spike-sorting method can be found in the Appendix.

When recording from multi-electrode arrays, each electrode records the signals from nearby neurons, and these signals can be combined across electrodes. Thus, by combining tetrodes together, we can gain additional spatial information about each signal. In this case, assuming that conditional on the spiking history and the current state variable the spiking activity is independent between tetrodes, we can augment

the observation distribution in Equation (3.5) to be

$$p(\Delta N_k, \vec{m}_k | x_k, H_k) \propto \begin{cases} \prod_{s=1}^S \exp[-\Delta_k \Lambda^s(t_k)], & \Delta N_k = 0; \\ \prod_{s=1}^S \prod_{i=1}^{\Delta N_k^s} [\lambda^s(t_k, \vec{m}_{k_i}) \Delta_k] \exp[-\Delta_k \Lambda^s(t_k)], & \Delta N_k > 0, \end{cases} \quad (3.13)$$

where S is the number of groups of recordings, in this case the number of tetrodes. ΔN_k^s is the number of spikes observed from neurons on tetrode s during $(t_{k-1}, t_k]$. $\Lambda^s(t_k) = \int_{\mathcal{M}} \lambda^s(t_k, \vec{m}_k) d\vec{m}_k$, where $\lambda^s(t_k, \vec{m}_k)$ defines the joint mark intensity function for neurons on tetrode s , where $s = 1, \dots, S$.

In this analysis, we propose a marked Poisson, nonparametric clusterless encoding model to estimate the joint mark intensity function $\lambda(t, \vec{m})$, given an observation interval $t \in (0, T]$ with a sequence of N spike times $0 < u_1 < \dots < u_i < \dots < u_N \leq T$:

$$\hat{\lambda}(t, \vec{m}) = \frac{\frac{1}{N} \sum_{i=1}^N \frac{1}{B_x} \frac{1}{\det(B_{\vec{m}})} \mathcal{K} \left(\frac{x(t) - x_{u_i}}{B_x}, B_{\vec{m}}^{-1} (\vec{m} - \vec{m}_{u_i}) \right)}{\frac{1}{T} \sum_{j=1}^T \frac{1}{b_x} K \left(\frac{x(t) - x_j}{b_x} \right)}, \quad (3.14)$$

where N is the total number of spikes, u_i is the time of i^{th} spike, and T is the total time of the experiment. \mathcal{K} is a multivariate kernel in both the place field and the mark space whose smoothness depends on both the smoothing parameter B_x and the bandwidth matrix $B_{\vec{m}}$ (Ramsay and Silverman, 2010). K is a univariate kernel on the spatial component only and its smoothing parameter is b_x .

Here we used Gaussian kernels for both \mathcal{K} and K :

$$\mathcal{K} = \left\{ \frac{1}{\sqrt{2\pi}} \exp \left[-\frac{1}{2} \left(\frac{x(t) - x_{u_i}}{B_x} \right)^2 \right] \right\} \cdot \left\{ \frac{1}{(\sqrt{2\pi})^d} \exp \left[-\frac{1}{2} (\vec{m} - \vec{m}_{u_i})^T B_{\vec{m}}^{-T} B_{\vec{m}}^{-1} (\vec{m} - \vec{m}_{u_i}) \right] \right\}, \quad (3.15)$$

$$K = \left\{ \frac{1}{\sqrt{2\pi}} \exp \left[-\frac{1}{2} \left(\frac{x(t) - x_j}{b_x} \right)^2 \right] \right\}, \quad (3.16)$$

where $d = 4$ is the dimension of the mark space. The bandwidths were chosen based on previous knowledge of two things: variability in the waveform structure for a single neuron and prior information about spatial extent of a place field. For this analysis, b_x and B_x were set to 1.5% of the linear track length, and $B_{\bar{m}}$ was set to a 4-dimensional scalar matrix whose diagonal entries equal 20 mV. We found that our results remain consistent for a wide range of bandwidths.

We then compare the results to decoding based on sorted data. These data included a total of 11 active place cells across an 840 second recording session. To assess whether the models generalize well to unobserved data as well as to limit overfitting, both algorithms are carried out using a 5-fold cross-validation.

Figure 3·4 illustrates the decoding results from signals recorded from 5 tetrodes in the CA1 and CA2 regions of the hippocampus using the two different algorithms. Here we display the first 125 seconds of decoding results in both panels. Figure 3·4A are decoding results using the clusterless decoding algorithm. Figure 3·4B are decoding results where the decoder is applied after the spikes have been manually sorted into clusters. The blue line shows the true position of the animal, and the red region shows the estimated posterior density at that time step. We can see that the estimate tracks the true position closely, and the region of highest posterior density covers the true position most of the time. In contrast, as shown in Figure 3·4B, the posterior density using decoding with spike-sorting has some trouble tracking the true trajectory, for example between time 6250 seconds and 6260 seconds.

To assess quality of fit, we calculated the root-mean-square error (rMSE) in centimeters and the percentage of the time that the true position values fall within the 99% HPD region. We compared the rMSE and percentage of time within 99% HPD region for the decoding results on the entire 840 seconds of recording. The posterior




Figure 3-4: Decoding results for hippocampus data. (A) Posterior density of the rat’s linear position using the clusterless decoding algorithm. The blue line represents the actual position of the rat. The red represents the posterior density at each time step. The posterior density computed through the clusterless decoding algorithm has an rMSE of 14.3 cm and remains within the 99% HPD region 74.25% of the time. (B) Posterior density of the rat’s linear position using decoding with manual spike-sorting. The posterior density computed by first spike-sorting and then decoding has a rMSE of 26.0 cm and remains within the 99% HPD region 70.76% of the time.

density computed through the clusterless decoding algorithm has an rMSE of 14.3 cm and remains within the 99% HPD region 74.25% of the time. The size of the 99% HPD region for the clusterless decoding algorithm has a mean of 10.5 cm and a standard deviation of 6.96 cm. The posterior density computed by first spike-sorting and then decoding has a rMSE of 26.0 cm and remains within the 99% HPD region 70.76% of the time. The size of the 99% HPD region for the decoded with sorted spikes has a mean of 17.2 cm and a standard deviation of 13.14 cm.

The clusterless algorithm has narrower 99% HPD region, more accurately reflecting the uncertainty of the decoding estimates. The algorithm that uses spike-sorting has wider 99% HPD region, suggesting less certainty in the estimates. Therefore the estimated certainty for the clusterless decoding method is both higher and more ac-

curate than the estimated certainty for the spike-sorting based method. Both visually and numerically, we showed that our proposed clusterless algorithm performs as well as or better than algorithm based on sorted spikes.

3.5 Discussion

We previously used point process theory to develop efficient decoding algorithms based on spike train observations (Brown et al., 1998; Eden et al., 2004; Huang et al., 2009; Koyama et al., 2010). In each of these cases, a key assumption is that the signals have been accurately sorted into single units before the decoding algorithm is applied. Although new spike-sorting algorithms are actively being developed, spike-sorting still remains a time-consuming, difficult task, suffers from many sources of errors, and likely provides biased estimates (Lewicki, 1998; Harris et al., 2000; Quiroga, 2012). In this article, we have proposed a novel clusterless decoding algorithm that maintains the accuracy of previous methods, but avoids spike-sorting.

The proposed clusterless decoding algorithm has several important advantages. First, this algorithm does not require that the multiunit activity has been accurately sorted into single units. Instead, by using the theory of marked point processes, the algorithm characterizes directly the relationship between a desired variable and features of the spike waveforms. We bypass spike-sorting by modeling the spiking activity as a joint function of the state variable to decode and of features of the spike waveforms. Therefore, this new algorithm can incorporate information from spikes that, in previous decoders that relied on spike sorting, may have been thrown out because of difficulties in clustering, or misclassified. Indeed, in a recent analysis of the effects of spike-sorting schemes on decoding performance, Todorova and colleagues confirmed that discarding waveforms that do not match any template—the “hash”—systematically degrades decoding (Todorova et al., 2014), consistent with previous

studies (Stark and Abeles, 2007; Fraser et al., 2009).

Another advantage of the marked point process approach is that the joint intensity mark function defines a population level place receptive field structure, which will typically have multimodal structure. That is, the joint intensity mark function is expected to have multiple peaks in the joint place and mark space. Previous work by Huang et al. has shown that decoding algorithms that allow for highly non-Gaussian and multimodal posterior densities perform better at reconstructing the animal's trajectory and predicting future decisions (Huang et al., 2009).

Last, formulating the decoding problem within a state-space paradigm offers a number of specific benefits. It allows us to incorporate knowledge about the system we are decoding, in this case the rat's position trajectory. This also imposes an implicit continuity constraint, preventing large fluctuations in the state estimates. This allows us to track the state variable smoothly without high dependence on the choice of time step during decoding. Moreover, due to the state-space approach, the resulting decoding algorithm incorporates information from both spike intervals and non-spike intervals.

We tested the new decoding algorithm on tetrode recordings from the hippocampus of a rat running back and forth on a linear track. We showed that the clusterless decoding algorithm provided a slightly improved accuracy than that of the decoding with manual spike-sorting approach. The intention of this example is to demonstrate comparable decoding accuracy without the additional, time-consuming spike-sorting step, thus making this algorithm a suitable candidate for real-time application.

There are a number of directions in which this work may be extended. The numerical integration we used to compute the filtering is efficient when decoding low-dimensional signals. When the desired variable is high-dimensional, alternative adaptive algorithms can be used to ensure efficiency, such as sequential Monte Carlo

methods or Gaussian approximate filters.

In the application example, we presented a nonparametric encoding model to estimate the joint mark intensity function based on kernel methods. We have found that our results remain consistent for a wide range of bandwidths. Therefore, in real-time decoding, it is reasonable to set the smoothing parameter and bandwidth matrix beforehand. We also recognize for non-parametric kernel based encoding, the computation time increases with the total number of spikes. In the interest of real-time applications, it is possible to explore other more efficient models as well. Within the family of nonparametric encoding model, one can include a time dimension with receding horizon to reduce the number of spikes encoded at any given time.

Another choice of encoding model is semi-parametric models, for example, a multi-dimension grid interpolated for the joint distribution of mark and place field. This grid can serve as a kind of look-up table for the estimated joint mark intensity function. By calculating this grid beforehand, we can perform real-time clusterless decoding without any computation time spent in encoding. A more general approach to reduce encoding time is to perform dimension reduction on the multi-dimensional mark space before encoding.

Finally, as we move away from the prerequisite of spike-sorting, multivariate marked point process models can be developed to describe coupling between neurons (Ba et al., 2014).

Another future role for these methods is in the development of new types of closed-loop experiments. Traditionally, experiments designed to study the role of specific spike patterns in stimulus-response tasks take one of two forms: observational studies that characterize statistical properties of neural activity during such tasks or interventional studies that broadly alter neural activity over an entire neural population or brain region. However, new closed-loop experiments designed based on the content

of neural signals aim to characterize causal relationships between neural activity and the biological and behavioral signals they encode. The proposed algorithm can allow investigators to manipulate millisecond-timescale spike patterns in a content-specific way, altering spiking activity related to certain neural patterns while leaving activity related to other patterns intact.

In conclusion, we develop a novel method for modeling neural response properties and decoding biological and behavioral signals by expanding the class of neural models to incorporate marked point processes. We found that the resulting decoding signals were estimated with higher accuracy and more confidence than traditional spike-sorting based methods. We believe this work has broad implications, allowing for better neural coding models across a wide range of brain areas and neural systems.

3.6 Appendix

3.6.1 Hippocampal data collection and preprocessing

The hippocampal data in this paper are from a single male Long-Evans rat (500-600 g) trained to alternate in a linear track for liquid reward (condensed milk). A microdrive array containing 21 independently movable tetrodes was implanted targeting the hippocampal cell layers according to University of California San Francisco Institutional Animal Care and Use Committee and US National Institutes of Health guidelines. All neural signals were recorded relative to a reference tetrode in the corpus callosum. Following data collection, electrode locations were verified histologically to localize to the CA1–CA2 region of the hippocampus.

Data were collected using the NSpike data acquisition system (L.M.F. and J. MacArthur, Harvard Instrumentation Design Laboratory). An infrared diode array with a large and a small cluster of diodes was attached to the preamps during recording. Following recording, the rats position on the track was reconstructed using a

semi-automated analysis of digital video of the experiment. Individual units (putative single neurons) were identified by clustering spikes using peak amplitude, principal components, and spike width as variables (MatClust, M. Karlsson) (Karlsson and Frank, 2008).

3.6.2 Gaussian approximation to the posterior density

Because in this paper the state variable is low-dimensional, we used a simple Riemann sum to perform numerical integration over the state space in order to solve the one-step prediction density in Equation (3.4). If the state variable is high-dimensional and the posterior density can be assumed to be approximately Gaussian, a stochastic state point process filter (SSPPF) can be used. The SSPPF is constructed using a Gaussian approximation to the posterior density. Detailed derivation of approximate Gaussian filters for temporal point processes can be found in (Brown et al., 1998; Eden et al., 2004). Here we provide an extension of the SSPPF to the marked case.

For estimating higher-dimensional state variables, a linear recursive Gaussian approximation to the posterior density at each time step, $p(x_k | \Delta N_k, m_k)$ can be constructed with the estimated conditional mean $x_{k|k}$ and variance $W_{k|k}$ as follows:

$$W_{k|k}^{-1} = (F_k W_{k-1|k-1} F_k^T + Q_k)^{-1} + \left[\frac{\partial^2 \Lambda_k}{\partial x_k \partial x_k^T} \Delta_k - \sum_{i=1}^{\Delta N_k} \frac{\partial^2 \log \lambda_k^i}{\partial x_k \partial x_k^T} \right]_{x_k = F_k x_{k-1|k-1}} \quad (3.17)$$

$$x_{k|k} = F_k x_{k-1|k-1} + W_{k|k} \left[\sum_{i=1}^{\Delta N_k} \left(\frac{\partial \log \lambda_k^i}{\partial x_k} \right)^T - \left(\frac{\partial \Lambda_k}{\partial x_k} \right)^T \Delta_k \right]_{x_k = F_k x_{k-1|k-1}}, \quad (3.18)$$

where, for simplicity, λ_k^i denotes $\lambda(t_k, \vec{m}_{k_i} | H_k)$ and Λ_k denotes $\Lambda(t_k | H_k)$
 $= \int_{\mathcal{M}} \lambda(t_k, \vec{m}_k | H_k) d\vec{m}_k$. In the posterior mean equation, the innovation includes two terms. The first term, $\frac{\partial \log \lambda_k^i}{\partial x_k}$, depends on the marks of the observed spikes. The second term $\frac{\partial \Lambda_k}{\partial x_k}$ includes Λ_k , which is the intensity function of the ground process

and relates to the intensity of observing any spikes regardless of the mark.

3.6.3 Sequential Monte Carlo decoding algorithm

When the state variable is high-dimensional and the posterior density cannot be assumed to be approximately Gaussian, another computationally efficient alternative is a sequential Monte Carlo algorithm. Point process adaptive filters using sequential Monte Carlo approximations to the posterior density have been developed in previous literature (Ergun et al., 2007; Meng et al., 2011). Here we provide a pseudo-code description of the algorithm with extension to marked point processes. This is a bootstrap filter, so the proposal distribution is based on the one-step prediction density from the previous time step.

At each time step t , the algorithm produces a collection of weighed samples, or particles, each containing proposed values for the state variable x_t . We construct estimates for the state variable by computing their sample means over all the particles, and construct approximate 95% confidence intervals.

1. Initialization:

Set $t = 0$ and for $i = 1, \dots, n$ particles, draw the initial states and parameters from an initial probability distribution and set the importance weight of the i^{th} particle $w_0^i = n^{-1}$ for all i . Set $t=1$.

2. Importance sampling:

Using the particles from the previous step which represent the one-step prediction density defined in equation (3.4) as the sampling distribution, update all of the states x_t .

Evaluate the importance weight of the i^{th} particle

$$w_t^i = w_{t-1}^i p(\Delta N_t, \vec{m}_k | H_t),$$

where $p(\Delta N_t, \vec{m}_k | H_t)$ is computed by equation (4.5) or (4.13).

Normalize the importance weights

$$\tilde{w}_t^i = \frac{w_t^i}{\sum_j w_t^j}.$$

3. Resampling:

Resampling can be performed at any fixed interval. Draw n particles $\{\tilde{x}_t^i : i = 1, \dots, n\}$ from $\{s_t^i : i, \dots, n\}$ using the residual resampling approach. Reset the weights to $\tilde{w}_t^i = n^{-1}$ to obtain the Monte Carlo estimate of the probability density

$$p(\tilde{x}_t | N_t) \approx n^{-1} \sum_{i=1}^n \delta(\tilde{x}_t - \tilde{x}_t^i),$$

where $\delta(\cdot)$ is the Dirac delta function, indicating a point mass at 0.

4. Repeat steps 2–3.

Chapter 4

Rapid Classification of Hippocampal Replay Content for Real-time Applications

4.1 Introduction

Because of the profound role memory plays in our lives, the study of the neural mechanisms that underlie memory processes, such as memory encoding, consolidation and retrieval, has attracted substantial interest among neuroscientists. Memory is thought to depend on the reactivation of patterns of activity related to previous experience, and efforts to identifying these patterns have focused, in part, on hippocampal sharp-wave ripple (SWR) events (Carr et al., 2011; Buzsaki, 2015). During these events, the hippocampus often replays sequences that represent previously experience behavioral trajectories (Redish and Touretzky, 1998; Redish, 1999; Lee and Wilson, 2002; Ji and Wilson, 2007; Foster and Wilson, 2006; Diba and Buzsaki, 2007). Sleep SWRs were originally proposed to support memory consolidation and related processes (Buzsaki, 1989) and during sleep, disrupting SWRs has been shown to lead to subsequent performance deficits in a spatial memory task (Girardeau et al., 2009; Ego-Stengel and Wilson, 2010).

Interestingly, replay during SWRs also occurs in the awake state, most prevalently during periods of slow movement and immobility (Foster and Wilson, 2006; Diba and Buzsaki, 2007; Csicsvari et al., 2007; Cheng and Frank, 2008; Karlsson and Frank,

2009; Singer and Frank, 2009; Pfeiffer and Foster, 2013; Wu and Foster, 2014; Silva et al., 2015). Interruption of these awake SWRs leads to a specific learning and performance deficit (Jadhav et al., 2012). That result helped establish the importance of SWRs, but does not establish that the specific pattern of spiking activity within SWRs contributes to memory processes. Individual SWRs can contain spiking patterns that replay different sequences, and further, these sequences can be replayed both in a temporally forward and reverse order during awake SWRs (Foster and Wilson, 2006; Csicsvari et al., 2007; Diba and Buzsaki, 2007; Wu and Foster, 2014; Ambrose et al., 2015). In particular, reverse replay is hypothesized to be an integral mechanism for learning about recent events and insights into this type of replay is could be critical to understanding how animals learn from experience (Foster and Wilson, 2006; Pfeiffer and Foster, 2013).

Therefore, a deeper examination of how replayed information contributes to learning and decision-making requires the ability to manipulate SWR events based on their content. This presents a fundamental analysis challenge: to decode and characterize replay content in real-time.

However, studies of SWR events have so far been mostly, if not all, limited to analyses of spiking data that were sorted off-line. Currently, real-time spike-sorting algorithms tend not to be sufficiently accurate to allow for closed-loop interventions. Additionally, methods for categorizing replay content typically require waiting for the completion of a replay event, for example, to compute a correlation measure to study temporal replay order (Foster and Wilson, 2006; Wu and Foster, 2014). The lack of dynamic estimation methods makes it difficult to further probe any causal relationship between a specific replay sequence and memory function that may happen downstream of the replay event. Among the currently used methods, there is no consensus as to which classifies the replay content most accurately and reliably.

One approach that has been successful in decoding dynamic content from neural data is state-space modeling (Brown et al., 1998; Smith and Brown, 2003; Eden et al., 2004; Truccolo et al., 2005; Srinivasan et al., 2006; Kemere et al., 2008; Wu et al., 2009; Huang et al., 2009; Paninski et al., 2010; Koyama et al., 2010). Here we incorporate state-space methods to develop a dynamic algorithm to determine if the content of a replay event represents a specific sequence, which will then make it possible to interrupt events based on the type of spatial trajectory they represent.

Previously, we and others have developed decoding methods that do not require multiunit spiking waveforms to be sorted into single units based on the theory of marked point processes (Chen et al., 2012; Kloosterman et al., 2014a; Deng et al., 2015). Here we extend the state of our marked point process filter to include a discrete state variable that categorizes the content of hippocampal replays in real time. By including such a “decision state”, we make use of the underlying state process in a content-specific way by selectively including different sources of information. In particular, it allows us to incorporate into our model three fundamental features of the information content of individual replay events: the location where the trajectory begins, whether the sequence of locations represents a trajectory proceeding toward or away from the animal’s current location, and whether the spiking pattern reflects place-field structure for a specific direction of movement. We then use the discrete decision state filter to categorize the replay content and compute our confidence about the classification. We illustrate our approach by decoding experimental data recorded in the hippocampus of a rat performing a spatial memory task.

4.2 Algorithm Development

Previous studies have used state-space models to track the actual position of a rat using spiking activity in hippocampus (Brown et al., 1998; Huang et al., 2009; Koyama

et al., 2010). In those studies, spiking activity was first modeled as a point process, which can be fully characterized by its conditional intensity function (Daley and Vere-Jones, 2003). The conditional intensity models then relates the firing activity to a state process that represents the animal’s actual position.

In this section, we generalize this method in two ways. First, we extend the conditional intensity function to allow for multi-unit activity. Second, we relate the firing activity to a new state process that represents features of a replay event rather than an actual movement trajectory. Here we posit that the dynamics of the replayed trajectory are similar to those of the actual movement, although with an accelerated time course. For this reason, we will call this new state, the “replay position.” We might loosely interpret this new state process as a kind of mental exploration or “mental time-travel” (Tulving and Markowitsch, 1998; Hasselmo, 2009). However, it is important to note that this new state does not necessarily represent any real or imagined position.

4.2.1 Discrete state point process filters

Any point process representing neural spiking can be fully characterized by its conditional intensity function (Daley and Vere-Jones, 2003). A conditional intensity function describes the instantaneous probability of observing a spike, given previous spiking history. By relating the conditional intensity to specific biological and behavioral signals we can specify a spike train encoding model.

The conditional intensity also generalizes to the marked case, in which a random vector, termed a mark, is attached to each point (Cox and Isham, 1980). Here we use the mark to characterize features of the spike waveform. In the case of tetrode recordings, for example, the mark could be a length four vector of the maximum amplitudes on each of the four electrodes at every spike time.

A marked point process is completely defined by its joint mark intensity

$$\lambda(t, \vec{m}|H_t) = \lim_{\Delta \rightarrow 0} \frac{\Pr(\text{a spike with mark } \vec{m} \text{ in } (t, t + \Delta]|H_t)}{\Delta}, \quad (4.1)$$

where H_t is the history of the spiking activity up to time t . $\lambda(t, \vec{m}|H_t)$ represents a joint stochastic model for the marks as well as the arrivals of the point process. The joint mark intensity characterizes the instantaneous probability of observing a spike with mark \vec{m} at time t as a function of any factors that may influence spiking activity.

Traditionally, we posit that the spiking activity depends on some underlying internal state variable $x(t)$, such as an animal’s location in space, that varies across time. Here, in order to allow for real-time decision making based on the information content of replay events, we extend the state variable from a single continuous state $x(t)$ to a joint state variable $\{x(t), I\}$ where $x(t)$ remains an underlying dynamic, continuous state process and I is a discrete, fixed state variable. The discrete variable I is constant in time but influences the dynamics of the continuous state $x(t)$. For example, this discrete state could be a binary variable that describes whether a movement is forward or reverse replay events. In closed-loop experiment, we might make a decision to intervene based on an estimated value of this discrete state. For that reason, we will use the term “decision state” to represent this variable.

A multi-unit encoding model is constructed by writing the joint mark intensity as a function of all the variables that influence the neural spiking,

$$\lambda(t, \vec{m}|H_t) = g(x(t), I, \vec{m}|H_t). \quad (4.2)$$

Here g is the function that describes the instantaneous probability of observing a spike with a waveform that gives mark \vec{m} at time t based on the joint state value. This joint conditional intensity models spiking activity across a population of neurons.

The conditional intensity of observing any spikes in this population can be com-

puted by integrating the joint mark intensity $\lambda(t, \vec{m}|H_t)$ over the entire space of possible marks, \mathcal{M} ,

$$\Lambda(t|H_t) = \int_{\mathcal{M}} \lambda(t, \vec{m}|H_t) d\vec{m}. \quad (4.3)$$

$\Lambda(t|H_t)$ can be understood as the conventional conditional intensity of a temporal point process that includes the spiking of all the neurons in the population. For a marked point process, $\Lambda(t|H_t)$ is called the intensity of the ground process (Daley and Vere-Jones, 2003). The mark space \mathcal{M} can be of any dimension.

A joint state point process filter applies the theory of state-space adaptive filters (Haykin, 1996) to compute, at each time point, the posterior distribution of the joint state variable given observed marked point process. To derive the filter in discrete-time, we first partition the observation interval $[0, T]$ into $\{t_k : 0 \leq t_1 < \dots < t_N \leq T\}$ and let $\Delta_k = t_k - t_{k-1}$. The posterior density for the joint state variable $\{x(t), I\}$ can be derived simply using Bayes's rule,

$$p(x_k, I|\Delta N_k, \vec{m}_k, H_k) = \frac{p(x_k, I|H_k)p(\Delta N_k, \vec{m}_k|x_k, I, H_k)}{p(\Delta N_k, \vec{m}_k|H_k)}, \quad (4.4)$$

where $x_k = x(t_k)$ is the continuous state variable at time t_k , ΔN_k is the number of spikes observed in the interval $(t_{k-1}, t_k]$, and H_k is the history of spiking activity up to time t_{k-1} . \vec{m}_k represents a collection of mark vectors \vec{m}_{k_i} , $i = 1, \dots, \Delta N_k$, observed in the interval $(t_{k-1}, t_k]$.

The first term in the numerator on the right-hand side of Equation 4.4, $p(x_k, I|H_k)$, is called the one-step prediction density for the joint state variable, and is determined by the Chapman-Kolmogorov equation as

$$p(x_k, I|H_k) = \int p(x_k|x_{k-1}, I) \cdot p(x_{k-1}, I|\Delta N_{k-1}, \vec{m}_{k-1}, H_{k-1}) dx_{k-1}. \quad (4.5)$$

Equation 4.5 has two components: $p(x_k|x_{k-1}, I)$, which is given by a state transition

model under the Markovian assumption conditioned on the discrete decision state, and $p(x_{k-1}, I | \Delta N_{k-1}, \vec{m}_{k-1}, H_{k-1})$, which is the posterior density from the last iteration step. We multiply this probability density function with the posterior distribution of the joint state variable at the previous time step t_{k-1} , and numerically integrate the product over all possible values of the continuous state at the previous time, x_{k-1} . The resulting integral is the a one-step prediction density at the current time t_k .

The second term in the numerator of Equation 4.4, $p(\Delta N_k, \vec{m}_k | x_k, I, H_k)$, is the likelihood or observation distribution of the population spiking activity at the current time conditioned on the joint state and past spiking history, which can be fully characterized by the joint mark intensity $\lambda(t, \vec{m} | H_t)$:

$$p(\Delta N_k, \vec{m}_k | x_k, I, H_k) \propto \begin{cases} \exp[-\Delta_k \Lambda(t_k | H_k)], & \Delta N_k = 0; \\ \prod_{i=1}^{\Delta N_k} [\lambda(t_k, \vec{m}_{k_i} | H_k) \Delta_k] \\ \cdot \exp[-\Delta_k \Lambda(t_k | H_k)], & \Delta N_k > 0. \end{cases} \quad (4.6)$$

We can interpret Equation 4.6 by separating the product on the right hand side into two terms. The $\prod_{i=1}^{\Delta N_k} [\lambda(t_k, \vec{m}_{k_i} | H_k) \Delta_k]$ term characterizes the distribution of firing ΔN_k spikes, such that the mark value of the i^{th} spike in the interval $(t_{k-1}, t_k]$ is m_{k_i} , where $i = 1, \dots, \Delta N_k$. If no spike occurs, i.e., $\Delta N_k = 0$, this term equals 1. The $\exp[-\Delta_k \Lambda(t_k | H_k)]$ term gives the probability of not firing any other spikes in the observation interval, where $\Lambda(t_k | H_k) = \int_{\mathcal{M}} \lambda(t_k, \vec{m}_k | H_k) d\vec{m}_k$ as defined in Equation 4.3.

The observation distribution is then multiplied by the one-step prediction density to get the posterior of the joint state at the current time:

$$\begin{aligned} & p(x_k, I | \Delta N_k, \vec{m}_k, H_k) \\ & \propto p(\Delta N_k, \vec{m}_k | I, x_k, H_k) \int p(x_k | x_{k-1}, I) \cdot p(I, x_{k-1} | \Delta N_{k-1}, \vec{m}_{k-1}, H_{k-1}) dx_{k-1}. \end{aligned} \quad (4.7)$$

Note that we can drop the normalization term in the denominator of Equation 4.4, $p(\Delta N_k, \vec{m}_k | H_k)$, because it is not a function of $\{x_k, I\}$.

Finally, because our goal is to estimate the discrete decision state, we integrate out the continuous state $x(t)$ from the posterior of $\{x(t), I\}$ in Equation 4.7 to compute the posterior probability of the discrete decision state I :

$$\Pr(I | \Delta N_k, \vec{m}_k, H_k) \propto \int p(I, x_k | \Delta N_k, \vec{m}_k, H_k) dx_k. \quad (4.8)$$

Substituting Equation 4.7 into Equation 4.8 yields a recursive expression for the estimation of the discrete decision state variable I :

$$\begin{aligned} & \Pr(I | \Delta N_k, \vec{m}_k, H_k) \\ & \propto \int p(\Delta N_k, \vec{m}_k | x_k, I, H_k) \int p(x_k | x_{k-1}, I) \cdot p(x_{k-1}, I | \Delta N_{k-1}, \vec{m}_{k-1}, H_{k-1}) dx_{k-1} dx_k. \end{aligned} \quad (4.9)$$

4.2.2 Algorithm discussion

The solution we obtained for the joint state filter in Equation 4.7 shares a similar form with solutions derived for a marked point process filter with only a continuous state $x(t)$ (Deng et al., 2015):

$$\begin{aligned} & p(x_k | \Delta N_k, \vec{m}_k, H_k) \\ & \propto p(\Delta N_k, \vec{m}_k | x_k, H_{k-1}) \int p(x_k | x_{k-1}) \cdot p(x_{k-1} | \Delta N_{k-1}, \vec{m}_{k-1}, H_k) dx_{k-1}. \end{aligned} \quad (4.10)$$

This similarity allows for an alternative interpretation of the new decision state filtering algorithm. For each possible value that the decision state, I , can take on, Equation 4.7 represents a separate filtering algorithm. We can therefore interpret Equation 4.7 instead as a system of equations describing d marked point process filters running in parallel, where d is the total number of values the decision state I

can take on. The value of the posterior integrated over x for each category relative to each of the others can be thought of as a weight, which describes the posterior probability for a particular value of the decision state. A pseudo-code description of this interpretation of the decision state filter is provided below. This alternative interpretation of the decision state filtering algorithm as weighted parallel continuous state filters in Equation 4.11b offers us a statistically intuitive connection between our proposed decision state filter and the classical state space filter.

From Equation set 4.11, we can see that the decision state filter obtains information from the following three distinct components:

$$\text{Initial condition: } p(I, x_0) = p(x_0|I) \Pr(I);$$

$$\text{State transition model: } p(x_k|x_{k-1}, I);$$

$$\text{Observation model: } p_{I=i}(\Delta N_k, \vec{m}_k|x_k, H_k),$$

$$i = 1, \dots, d.$$

For example, we might let the continuous state, $x(t)$, be the replay position and the discrete state, I , a binary indicator where $I = 0$ defines an outbound replay trajectory from point A to point B and $I = 1$ defines an inbound replay trajectory from point B to point A. Correspondingly, the three sources of information can be interpreted as

Initial condition: Where the replay trajectory starts;

State transition model: Whether the replay path follows
an outbound or inbound movement trajectory;

Observation model: Whether the replay activity reflects
spiking from inbound or outbound movements.

Pseudo-code description of decision state filter

1. Initialize the joint state variable $\{x(t), I\}$:

$$p(x_0, I) = p(x_0|I) \Pr(I) \quad (4.11a)$$

2. For time step t_k , we first solve multiple marked point process filters $p_{I=i}(x_k|\Delta N_k, \vec{m}_k, H_k)$ in parallel, where $i = 1, \dots, d$, and d is the total number of values the decision state I can take on:

$$p(x_k, I|\Delta N_k, \vec{m}_k, H_k) = \begin{cases} p_{I=1}(x_k|\Delta N_k, \vec{m}_k, H_k); \\ p_{I=2}(x_k|\Delta N_k, \vec{m}_k, H_k); \\ \dots\dots\dots \\ p_{I=D}(x_k|\Delta N_k, \vec{m}_k, H_k). \end{cases} \propto \begin{cases} p_{I=1}(\Delta N_k, \vec{m}_k|x_k, H_k) \int p_{I=1}(x_k|x_{k-1}) \\ \quad \cdot p_{I=1}(x_{k-1}|\Delta N_{k-1}, \vec{m}_{k-1}, H_{k-1}) dx_{k-1}; \\ p_{I=2}(\Delta N_k, \vec{m}_k|x_k, H_k) \int p_{I=2}(x_k|x_{k-1}) \\ \quad \cdot p_{I=2}(x_{k-1}|\Delta N_{k-1}, \vec{m}_{k-1}, H_{k-1}) dx_{k-1}; \\ \dots\dots\dots \\ p_{I=d}(\Delta N_k, \vec{m}_k|x_k, H_k) \int p_{I=d}(x_k|x_{k-1}) \\ \quad \cdot p_{I=d}(x_{k-1}|\Delta N_{k-1}, \vec{m}_{k-1}, H_{k-1}) dx_{k-1}. \end{cases} \quad (4.11b)$$

3. Each parallel filter is then weighted by the integrated value of the posterior computed for its category relative to others:

$$\Pr(I = j|\Delta N_k, \vec{m}_k, H_k) = \frac{\int p_{I=j}(x_k|\Delta N_k, \vec{m}_k, H_k) dx_k}{\sum_{i=1}^D \int p_{I=i}(x_k|\Delta N_k, \vec{m}_k, H_k) dx_k} \quad (4.11c)$$

Thus, to estimate the decision state, we must define probability models for $p(x_0|I)$, $p(x_k|x_{k-1}, I)$, and $p(\Delta N_k, \vec{m}_k|I, x_k, H_k)$.

4.3 Data Analyses

To demonstrate the ability of our proposed algorithm to estimate and classify replay content, we applied it to the analysis of an example set of electrophysiological data from a rat performing an alternation task in an M-shaped maze. We begin by defining probability models to capture the three sources of information discussed above that contribute to the decoding of replay trajectory. We then show the application of our marked point process filter using only a continuous state variable to decode the full trajectory of four example replay events. Next, we extend the state variable

to jointly include a discrete decision state that identifies whether each replay event represents an outbound or inbound trajectory as well as whether the temporal order of activity occurs forward or backward in time. We categorize the content of the four example replay events using the marked point process filter with the joint state variable. Finally, we investigate how thresholding the posterior probability of the decision state at different values affects the classification results during a recording epoch with multiple replay events.

4.3.1 Description of the experimental data

The experimental methods are described in detail in (Karlsson and Frank, 2008). In brief, data was taken from one male Long Evans rat implanted with a microdrive array containing 30 independently movable tetrodes targeting CA1 and CA3. Tetrodes that never yielded clusterable units across the 8 days of experiments were excluded from analysis. Spiking activity recorded on a total of 18 tetrodes was used for the data analyses below.

Ripples were detected on 11 of the 30 tetrodes recorded from CA1 by using an aggregated measure of the root mean square (RMS) power in the 150–250Hz band across the tetrodes (Csicsvari et al., 1999). The aggregated RMS power was then smoothed with a Gaussian (4 ms standard deviation) and SWR events were detected as lasting at least 15 ms above 2 standard deviation of the mean. The entire SWR time was then set to include times immediately before and after the power exceed the mean.

The data were recorded while the rat ran a continuous alternation task on an M-maze (76 cm \times 76 cm with 7 cm wide track) shown in Figure 4.1A. The animal was rewarded each time he visited the end of an arm in the correct sequence, starting in the center and then alternating visits to each outer arm and returning to the center. The animal would start at the food well in position O and run toward the

intersection, or “choice-point” (position CP), at the top of the center stem, where a choice would need to be made. The correct choice is to alternate between left and right on successive turns. If a correct choice is made, for example, to turn right, the animal would continue to run toward the food well at position R to receive a reward, and then return to the center well at position O to move on to the next trial. If an error is made, the rat is not given a reward and must return to the center well to initiate the next trial. All animal procedures and surgery were in accordance with University of California San Francisco Institutional Animal Care and Use Committee and US National Institute Health guidelines.

We linearized the actual 2D coordinate position of the rat to a single coordinate. The 1D coordinate indicates the total distance from the center well (position O) in centimeters, with negative numbers indicating trajectories that include a left turn and positive numbers indicating ones that include a right turn. When the rat was on the center arm of the maze, the region to which its position was mapped was determined by the direction from which the rat came during inbound trajectories and by the direction it would turn next when it reached the choice point (position CP) during outbound trajectories. Throughout this paper, to facilitate ease of visualization, when plotting, we label the linearized maze at position O (“origin”), CP (“choice-point”), L (“left food-well”), and R (“right food-well”) instead of the corresponding signed 1D coordinate.

4.3.2 Analysis of individual replay events: sources of information for decoding replay position

We first estimate a single state variable representing replay trajectory from unsorted ensemble spiking activity by using a marked point process filter. Because we are interested in decoding the replay trajectory, the state variable is simply the replay position, $x(t)$. Figure 4-1B–D illustrate the three sources of information that contribute to the

decode of replay position.

Figure 4-1B plots the initial condition of the state variable, $p(x_0)$. Here we assume a replay event is equally likely to start at any location on the maze by defining a uniform probability model for x_0 . Figure 4-1C plots the state transition probability matrix $p(x_k|x_{k-1})$ of transition probabilities from every possible replay position at the prior time step to each value at current time. Here we posit that the dynamics of the replayed trajectory are similar to those of the actual movement and use the empirical movement data to compute the state transition model. The empirical state transition matrix has most of its probability mass near the diagonal, reflecting continuous trajectories in each of the track section, with a few off-diagonal elements that reflect the presence of the choice point.

Figure 4-1D plots the estimated observation model for a single example tetrode. It characterizes the distribution of the spiking activity as a function of the replay position and features of the spike waveform (peak height of each channel), which is determined by the joint conditional intensity $\lambda(t, \vec{m}|H_t)$. In this case, we use a Poisson, nonparametric, kernel-based encoding model to estimate the joint conditional intensity (Deng et al., 2015). We describe the construction of bootstrap confidence bounds on $\lambda(t, \vec{m})$ in the Appendix, which takes into account the uncertainty in the estimation of the joint conditional intensity.

Each panel in Figure 4-1D plots a 2-dimensional slice of the 5-dimension position-mark model. We observe that this example tetrode picked up a lot of low amplitude spiking activity with peaks between -50 mV and 100 mV. There is also evidence for the presence of multiple place-cells. We observe, for example, from the plot showing “Channel 1” and “Position”, that the ensemble spiking activity on this example tetrode includes clusters with peaks above 100 mV that are place-specific. The firing rate is higher when the rat is at the choice point CP on a future left turn than on a

future right turn. When the rat is on a left turn, the firing rate increases near the left food-well (position L) but not at the right corner of the maze. Contrarily, when the animal is on a right turn, the firing rate increases at the right corner but not near the right food-well (position R). For the decoding analysis described below, we fit such marked point process models for a total of 18 tetrodes.

4.3.3 Analysis of individual replay events: decoding replay position

Figure 4.2 shows the decoded trajectories of four example replay events using a marked point process filter whose state variable is the replay position. The schematic representation of the decoding results is illustrated by the M-shaped maze on top right corner of each panel in Figure 4.2A–D, where the yellow arrowed line shows the decoded replay trajectory schematically. Each panel on the left in Figure 4.2A–D plots the unsorted ensemble spiking activity on each tetrode as a function of replay time. Each panel on the right in Figure 4.2A–D plots the posterior density for the replay position of the animal during each replay event as a function of time, i.e. $p(x_k | \Delta N_k, \vec{m}_k, H_k)$ as computed from Equation 4.10. The yellow region in the heat plot represents position where the estimated posterior density of the replay position at each time step is large.

We can see that replay example 1, plotted in Figure 4.2A, shows a replay trajectory that begins at the center-well (position O), moves up the center arm to the choice point (position CP), then turns right, and proceeds to reach the right well (position R). Replay example 2, plotted in 2B, shows a replay trajectory that begins at the center-well (position O), moves up the center arm to the choice point (position CP), and then turns left towards the left corner. The decoding results for replay example 4, plotted in Figure 4.2D, shows a replay trajectory that originates from the left well (position L), moves up the left arm to the left corner and towards the choice-point (position CP), and then down to the center well (position O).

The decoding results for replay example 3, plotted in Figure 4.2C, shows a replay trajectory that is first an outbound trajectory but then shifts to an inbound trajectory: the replay trajectory first begins from the center well (position O) and moves up the center-stem to the choice point (position CP), but then at around 70 ms into the replay, it moves inbound back to the center well (position O). However, we notice that the outbound trajectory replayed during the first 70 ms has a lower posterior probability density at about 0.1 (lighter shade of yellow) while the rest of the replay has a higher posterior probability density at about 0.3 (darker shade of yellow).

4.3.4 Analysis of individual replay events: sources of information for decoding a decision state

To categorize the content of hippocampal replays, our model requires an explicitly-defined, fixed, discrete decision state I . Previous studies have reported the presence of both forward and reverse replay events in rat hippocampus (Foster and Wilson, 2006; Csicsvari et al., 2007; Diba and Buzsaki, 2007; Wu and Foster, 2014; Ambrose et al., 2015). We therefore define a decision state I that consists of two components: movement direction with regard to the center well (“outbound” or “inbound”) and the temporal replay order of the actual spatial experience (“forward” or “reverse”).

$$I = \begin{cases} 1, & \text{outbound, forward;} \\ 2, & \text{outbound, reverse;} \\ 3, & \text{inbound, forward;} \\ 4, & \text{inbound, reverse.} \end{cases}$$

An “outbound, forward” replay event can be interpreted as a spiking sequence that reflects spiking activity during an actual outbound movement being replayed in forward time. Loosely we might envision this event as a mental exploration of

the rat moving from the center-well to a side-arm with its head facing forward. An “outbound, reverse” replay event can be interpreted as a spiking sequence that reflects spiking activity during an actual outbound movement being replayed backward in time. We might envision this event as a movie of the same mental exploration of an “outbound, forward” event with the frames played backward in time. Watching such a movie, one might see the rat moving from a side-arm to the center-well with its head facing away from the center-well with each backward step.

An “inbound, forward” replay event can be interpreted as a spiking sequence that reflects spiking activity during an actual inbound movement being replayed in forward time. We might envision this event as a mental exploration of the rat moving from a side-arm to the center-well with its head facing forward. An “inbound, reverse” replay event can be interpreted as a spiking sequence that reflects spiking activity during an actual inbound movement being replayed backward in time. We might envision this event as a movie of the same mental exploration of an “inbound, forward” event with the frames played backward in time. Watching such a movie, one might see the rat moving from the center-well to a side-arm with its head facing towards from the center-well with each backward step.

Recall in the “Algorithm Discussion” subsection, we defined three sources of information for our decoding algorithm. The first two sources of information, the initial condition of the state variable and the state transition model, contribute to the decoder knowledge about the spatial evolution of the movement trajectory. The last source of information, the observation model, contributes to the decoder knowledge about mark-place preference in ensemble spiking. Therefore, if a replay event is categorized as outbound, forward, it means that both the spatial evolution of the replay movement and the spiking are consistent with those during an actual outbound movement. If a replay event is categorized as inbound, forward, it means that both the

spatial evolution of the replay movement and the spiking are consistent with those during an actual inbound movement. However, if a replay event is categorized as a reversely replayed outbound trajectory, or simply outbound, reverse, it means that although the spiking during the replay event is consistent with that during an actual outbound movement, the spatial evolution of the replay movement is more consistent with that of an inbound movement. Similarly, if a replay event is categorized as a reversely replayed inbound trajectory, or simply inbound, reverse, it means that although the spiking during the replay event is consistent with that during an actual inbound movement, the spatial evolution of the replay movement is more consistent with that of an outbound movement.

Figure 4·3 provides an example of the three different sources of information, as defined in Equation set 4.11 in the pseudo-code description of the decision state filter. As a visualization aid, we plot sources of information that are consistent with those during outbound movements in shades of blue and sources of information that are consistent with those during inbound movements in shades of red.

In each of the subplots A–D of Figure 4·3, the left panel plots $p(x_0|I)$, the probability model for the initial condition of the replay position state variable conditioned on the discrete decision state. When $I = 1$ and $I = 4$, $x_0|I$ is the initial condition for all outbound trajectories and we assume it to have a Gaussian distribution centered at the center-well (position O), which is plotted in blue. When $I = 2$ and $I = 3$, $x_0|I$ is the initial condition for all inbound trajectories. Here we assume it to have a probability distribution that is constant away from the center-well and subtracts a Gaussian density near the center-well, which is plotted in red.

The center panels of Figure 4·3 show a slice of the transition matrix of the replay position (when the animal is at a position on the left arm shown as the black dot) conditioned on the decision state, $p(x_k|x_{k-1}, I)$. When $I = 1$ and $I = 4$, $p(x_k|x_{k-1}, I)$

is the state transition model for all outbound trajectories, plotted in blue. When $I = 2$ and $I = 3$, $p(x_k|x_{k-1}, I)$ is the state transition model for all inbound trajectories, plotted in red. In this analysis, when specifying both the forward and reverse events, we use the empirical state transition matrix computed from actual movement data. Another choice of the state model, when defining reverse events, can be the inverse of the state transition matrix computed from actual movement data. More specifically, when $I = 2$ or in the case of a reversely replayed outbound trajectory, the state model is the inverse of the state transition matrix computed from actual outbound movement data. When $I = 4$ or in the case of a reversely replayed inbound trajectory, the state model is the inverse of the state transition matrix computed from actual inbound movement data. In this case, we found that our decoding results were similar using either set of state transition model.

Visually, the center panels Figure 4-3 look very similar, due to the fact that locally the state transition model is very flexible. By carefully comparing the center panels plotted in blue with those plotted in red, we can see that if the animal is in the middle of the left arm, it is more likely to move towards the left well (position L) during an outbound trajectory and more likely to move towards the center during an inbound trajectory.

The right panels of Figure 4-3 show a slice of the estimated joint mark intensity $\lambda_I(t, \vec{m}|H_k)$ marginalized over three tetrode channels, leaving the linearized position and one remaining channel, on an example tetrode. When $I = 1$ and $I = 2$, $\lambda_I(t, \vec{m}|H_k)$ is the observation model for outbound movements, plotted in a blue shade. When $I = 3$ and $I = 4$, $\lambda_I(t, \vec{m}|H_k)$ is the observation model for inbound movements, plotted in a red shade. By comparing the observation models for outbound (in blue) and inbound (in red) trajectories, we can see that the ensemble firing activity observed on this example tetrode shows directional preference. For example,

the outbound observation model (in blue) shows the ensemble is most likely to spike with a peak height around 220 mV when the rat is near the choice point of a future left path. The inbound observation model (in red) shows the ensemble is most likely to spike with a peak height around 240 mV when the rat is returning to the choice point from a right turn.

4.3.5 Analysis of individual replay events: decoding a decision state

Figure 4-4 shows the decoding results of the decision state for the same four example replay events shown in Figure 4-2. The schematic representation of the decoding results is illustrated by the M-shaped maze on the top right corner of each panel in Figure 4-4, where the yellow arrow represents the decoded replay trajectory. The black triangle represents the replayed orientation of the rat at the beginning of an event in order to ease the visualization, especially in the case of reverse replay.

Each main panel in Figure 4-4A–D plots the estimated posterior probability of each of the four categories defined in the decision state as a function of replay time. The line in darker blue represents $p(I = 1 | \Delta N_k, \vec{m}_k, H_k)$ as computed from Equation 4.9, the probability of the replay event being an “outbound, forward” trajectory; the line in lighter blue represents the probability of the replay event being a reversely replayed outbound trajectory; the line in darker red represents the probability of being an inbound and forward trajectory from the side-well; the line in lighter red represents the probability of being a reversely replayed inbound trajectory from the side-well.

Figure 4-4A shows that by 60 ms into replay example 1, the decision state decoder has determined the event to be an outbound, forward event with almost 100% certainty. Figure 4-4B shows that by 60 ms into replay example 2, the decoder has determined the event to be an inbound, reverse event with near 100% certainty. Figure 4-4C shows a decoding trajectory that initially suggests an outbound, forward event

with high probability but after around 100 ms increasingly suggests an inbound, forward event. Figure 4.4D shows that by 25 ms into replay example 4, the decision state decoder has determined the event to be an outbound, reverse event with near 100% certainty. Near 100 ms, this confidence is briefly reduced, but quickly returns to almost 100% confidence.

For some of these examples, if we were to make a decision at the first time the posterior probability of the decision state passed a fixed confidence threshold, the resulting decision would change based the threshold level. For example, for replay 1, if we were to set a moderate confidence threshold of 0.8, we might have determined the replay to be an inbound, reverse event at around 25 ms into the ripple. Similarly, for replay example 3, even with a relatively high threshold of 0.9, we would have determined the replay to be an outbound, forward event. These observations of switching between directionality and temporal replay order within a replay event are similar to previous findings that support the notion that replay captured the unique structure of the environment (encoding of the choice point) (Wu and Foster, 2014).

As explained in the *Algorithm Discussion* subsection, the decision state decoder can be interpreted as a set of weighted, marked point process filters running in parallel. In Figure 4.5, we illustrate this alternative interpretation by plotting $p_{I=i}(x_k | \Delta N_k, \vec{m}_k, H_k)$, $i = 1, 2, 3, 4$ for each example replay event as computed from Equation 4.11c. The four panels in each subplot of Figure 4.5A–D shows the decoded replay trajectory conditioned on each of the four decision states.

For each replay example shown in Figure 4.5, when the posterior probability for a particular decision state is low, very little of the probability mass appears in the corresponding panel. For example, in Figure 4.5A, there is no visible probability mass in the outbound, reverse or inbound, forward panels, and only fleeting presence of probability mass in the inbound, reverse panel. After 50 ms, nearly all of the

probability mass is completely contained in the outbound, forward panel: the replay trajectory moves from the center-well to the choice-point, and then to the right well; this is consistent with the visualization in Figure 4.4A, which shows that this replay event is initially most consistent with an inbound, reverse movement during the first 50 ms, but is eventually estimated to be an outbound, forward movement with high confidence.

The probability mass of the estimated replay trajectory plotted in Figure 4.5B is primarily located in the inbound, reverse panel. The trajectory moves from the center-well to the choice-point, then to the right corner, a sequence of position similar to the trajectory in Figure 4.5A. However, the spiking sequence is more consistent with spiking during inbound movements, which gives us confidence that the trajectory in Figure 4.5B is actually a reverse event. Similarly, the probability mass of the replay trajectory plotted in Figure 4.5D is primarily located in the outbound, reverse panel. While it shows a trajectory that moves from the left well to the choice-point and then to the center-well, the spiking sequence is more consistent with spiking during outbound movements, which gives us confidence that it is actually a reverse event.

The estimated trajectory in Figure 4.5C has probability mass initially concentrated in the outbound, forward panel and then after 100 ms becomes concentrated in the inbound, forward panel. The replay trajectory first moves outbound from the center-well to the choice-point, and then moves inbound back to the center-well, consistent with the shift in decision state shown in Figure 4.4C. This result suggests our model, which assumes a constant decision state over an entire replay event, may be misspecified. We could augment the state model to allow for discrete transitions that occur in mixed replay events (Wu and Foster, 2014).

4.3.6 Analysis of an entire recording epoch: decoding a decision state

Because our algorithm was developed with the goal of real-time manipulation of SWR replays, we investigated the amount of time required to make a decision to interrupt a replay event based on our desire confidence in its classification. Figure 4-6 plots the histogram of the fraction of replay events that pass a specified confidence threshold on the estimated decision state as a function of time into a replay event for an entire recording epoch. Events that can be classified are plotted in green. Events whose decision state probabilities did not pass the threshold by the end of a replay and cannot be classified are plotted in black. With a moderate threshold of 0.8, less than 2% of replay events failed to be categorized into any of the four categories, and the majority of the events were categorized within 20 ms of the replay onset. As we increase the threshold value, which means requiring more confidence of the classification, more replays failed to be categorized, and we have to wait longer into a replay for it to be categorized.

Figure 4-7A and 5-7B plot the decoding results for the decision state of all replay events by lap using a threshold of 0.9 on the classification confidence for two entire recording epochs on the same experiment day. Figure 4-7A plots results when the rat explored a novel track while Figure 4-7B shows results from a familiar track. Events that are classified into “outbound, forward”, “outbound, reverse”, “inbound, forward”, and “inbound, reverse” are plotted in darker blue, lighter blue, darker red, and lighter red, respectively. Events whose decision state probabilities did not pass the threshold by the end of a replay and cannot be classified are plotted in black. Laps are indexed on the x-axis. Left and outbound laps, left and inbound laps, right and outbound laps, and right and inbound laps are denoted on the x-axis as “oL”, “iL”, “oR”, and “iR”, respectively.

We can see from Figure 4-7A and 5-7B that, for both example epochs, not only

does the number of replay events on each lap vary, so does the proportion of replay events that fall into each decision category. In Figure 4.7A, which plots the decoding results when the rat was exploring a novel track, the majority of the replays are inbound, reverse events colored in lighter red. In contrast, in Figure 4.7B, which plots the decoding results when the rat was running on a familiar track, the majority of the replays are either outbound, forward events colored in darker blue or outbound, reverse events colored in lighter blue, with only one inbound, reverse event colored in lighter red.

In this small data sample, there appears to be a difference in the frequency of outbound, forward versus inbound, reverse replay events between novel and familiar conditions. While it is too early to draw any conclusions about the distribution of replay events in novel versus familiar exposures to an environment, this example highlights the manner in which our decoding approach can be used to better understand the role of replay in learning and memory tasks.

4.4 Discussion and Conclusion

Previous work has shown that interruption of awake SWRs leads to a specific learning and performance deficit (Jadhav et al., 2012). However, outstanding questions about the role of specific replay sequences in learning and memory remain. Here we present an approach to efficiently extract information about the content of SWR replay events that can facilitate real-time experiments to manipulate neural activity in a content-specific manner. Such an experiment has the potential to improve our understanding of the causal role that replay has in learning and memory. We extend our previous work on marked point process filters to capture the tuning properties of multiunit activity without spike sorting (Deng et al., 2015). Our algorithms makes use of three potential sources of information to classify replay events.

We illustrated the properties of the algorithm on four example replay events recorded as a rat performed an navigational memory task. We defined a decision state that captured both the movement direction (“outbound” or “inbound”) and the temporal order (“forward” or “reverse”) of the spatial experience being replayed. We demonstrated that our decoder provided a rapid classification of the replay content and a measure of classification uncertainty.

Our classification algorithm combines information from three sources—where the replay begins, the spatial sequence of the trajectory, and with which movement direction the neural activity is consistent. This allows us to express our confidence about our classification estimates based on the consistency of information across these sources, without the need for template matching for each category individually. For example, replay example 1 (Figure 4-4A) showed a switch at about 60 ms into the replay from being classified as a reverse, inbound trajectory ($I = 4$) with about 80% confidence by the decoder to a forwardly replayed outbound trajectory ($I = 1$) with full certainty. These observations of switching between directionality and temporal replay order within a replay event are similar to previous findings of mixed replay content that depends on environmental features such as choice-points (Wu and Foster, 2014).

A closer look at the sources of information that are combined to estimate the decision state for replay example 1 (Figure 4-3A and 5-3D) reveals that the switch can be attributed to a change in the spike and waveform patterns. Throughout the replay event, the spatial evolution remains consistent with an outbound trajectory; both states $I = 1$ and $I = 4$ are consistent with an initial condition in the center-arm and a state transition model that moves away from the center-arm. However, $I = 4$ is defined by using an observation model that is consistent with actual inbound spatial experiences while $I = 1$ is defined by using an observation model that is consistent

with actual inbound spatial experiences. Therefore, our algorithm allows us both to estimate and classify replay trajectory rapidly, and to study these replay events in a more feature-specific way.

We also applied our algorithm to decode the decision state of two longer recording epochs, each at least 15 minutes in length: one when the rat was exploring a novel track and the other a familiar track. We showed that within these two example epochs, shown in Figure 4.7, there was a persistent pattern of composition of classes of replay events across laps. Between these two epochs, however, the distributions of classes of replay events are quite different. Specifically, we found that inbound, reverse events were more common in our example familiar epoch. There are a number of possible explanations for this difference in replay distribution. For example, Karlsson and Frank 2009 suggested that some replay events may represent previously experienced environment—a situation that may occur more frequently in familiar environments. In future work, we plan to use these methods to explore replay of local versus previously experienced environments.

There are a few additional caveats related to these methods and the example analyses above. We assumed that the joint-mark intensity model that we computed from the data was known with complete certainty. By incorporating the encoding model uncertainty, as described in the Appendix, we would expect the confidence level for each decoded replay event to be decreased. Additionally, our state model assumed that each replay event was one of four possible types. In fact, the discussion above highlights other possible origins for the observed replay sequences. Similarly, previous analyses have found evidence for mixed events (Wu and Foster, 2014), which does not match the model assumptions. The effect of model misspecification on these methods is an important direction for future research.

There are a number of directions in which this work may be extended. For exam-

ple, one question is whether there is a multiunit receptive field model better suited for real-time estimation than the kernel models we used here. We are actively investigating the potential for Gaussian mixture models in describing the ensemble spiking activity. Another extension is to explore more features of the spike waveform as marks within our model, such as including its width and the slope.

Closed-loop experiments are increasingly common in neuroscience. They allow for real-time control over neural dynamics and animal behavior with the goal of probing the causal relationship between neural activity and behavior and the testing of hypotheses of the underlying neural systems that would be difficult or impossible to address in an open-loop setting (Siegle and Wilson, 2014; Grosenick et al., 2015). Closing the loop demands novel data analysis methods that can guide informed experimental decisions—stimulation, inhibition, and modulation—using the neural activity on a physiologically relevant timescale.

This work provides an approach to decode continuous features and make discrete decisions based on the confidence gained from multiunit spiking activity with a short latency on the order of milliseconds, which is suitable for closed-loop control. Components of our proposed state-space filtering algorithm can be formalized to test specific hypothesis, such as the causal role of specific activity patterns in generating subsequent activity patterns and behavior. We envision the need for this algorithm to broaden and enhance our understanding of a wide range of neural systems.

4.5 Appendix

4.5.1 Bootstrap confidence bounds for the estimated joint-mark intensity

We can estimate the uncertainty associated with the joint-mark intensity model analytically or numerically. Here we describe a bootstrap algorithm for constructing confidence bounds on the joint-mark intensity. Over the encoding period, we resample

each discrete time step with replacement to generate a new set of positions, spikes, and marks. This set is the j th out of B total bootstrap samples. For each bootstrap sample, we construct a separate joint-mark intensity estimate, $\hat{\lambda}^{(j)}(t, \vec{m}|H_t)$. The point-wise confidence bounds can be computed using the quantiles of the collection of bootstrap estimates at any mark value.

4.5.2 Incorporating model uncertainty into decoding algorithm

In the analyses described in this paper, we assumed that the joint-mark intensity was estimated with perfect confidence. In order to incorporate the uncertainty in the joint-mark encoding model, we can augment the likelihood in Eq. 9 by using the bootstrap joint-mark intensity computed above. The likelihood now becomes:

$$\begin{aligned} p(\Delta N_k, \vec{m}_k | x_k, I, H_k) &= \int p(\Delta N_k, \vec{m}_k | x_k, I, H_k, \lambda) p(\lambda) d\lambda \\ &= \frac{1}{B} \sum_k \sum_j \hat{p}^{(j)}(\Delta N_k, \vec{m}_k | x_k, I, H_k), \end{aligned}$$

where $\hat{p}^{(j)}(\Delta N_k, \vec{m}_k | x_k, I, H_k)$ represents the probability, as in Eq. 6, computed using the j th bootstrap joint-mark intensity $\hat{\lambda}^{(j)}(t, \vec{m}|H_t)$.

The filtering algorithm then proceeds as previously described using this augmented likelihood.

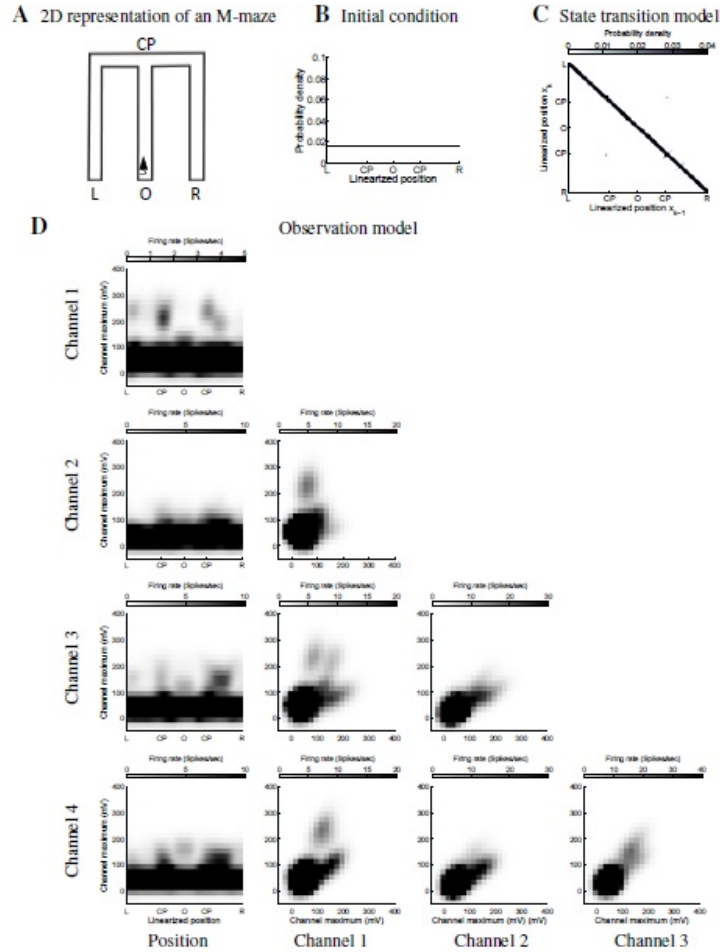


Figure 4-1: Linearization and encoding models for position. (A) 2D representation of an M-maze. The black triangle represents an animal whose head is oriented towards the choice-point. (B) Initial condition of the state variable, linearized position $x(t)$. Here we assume a replay event is equally likely to start at any position on the maze. (C) State transition probability matrix $p(x_k|x_{k-1})$ of transition probabilities from every possible replay position at the prior time step to each value at current time. Here we posit that the dynamics of the replayed trajectory are similar to those of the actual movement and use the empirical movement data to compute the state transition model. (D) Observation model for an example tetrode. Here we plot the pairwise marginalization of the 5-dimension position-mark model. Panel at the intersection of each row and column plots the estimated joint mark intensity function marginalized over the row and column linearized position or tetrode channel.

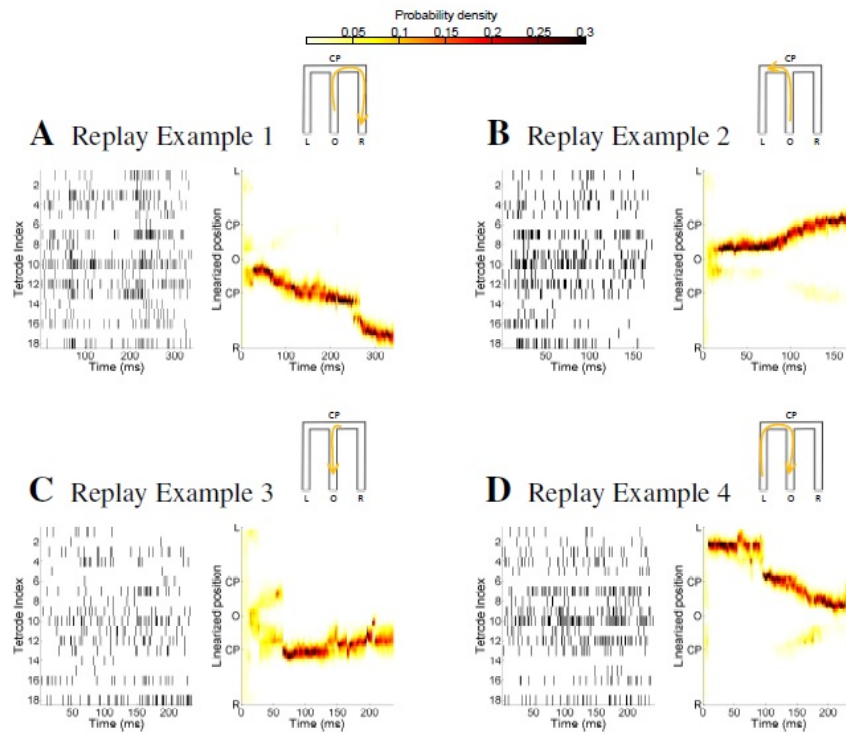


Figure 4.2: Decoding replay trajectory from multiunit spiking activity using a marked point process filter. (A, B, C, D): *Left* Unsorted ensemble spiking activity on each tetrode as a function of replay time. *Right* Replay trajectories of four examples decoded using clusterless methods where the heat plot shows the estimated posterior density at each time step. *Top* Schematic representation of the decoding results where the yellow arrowed line illustrates the evolution of the replay trajectory.

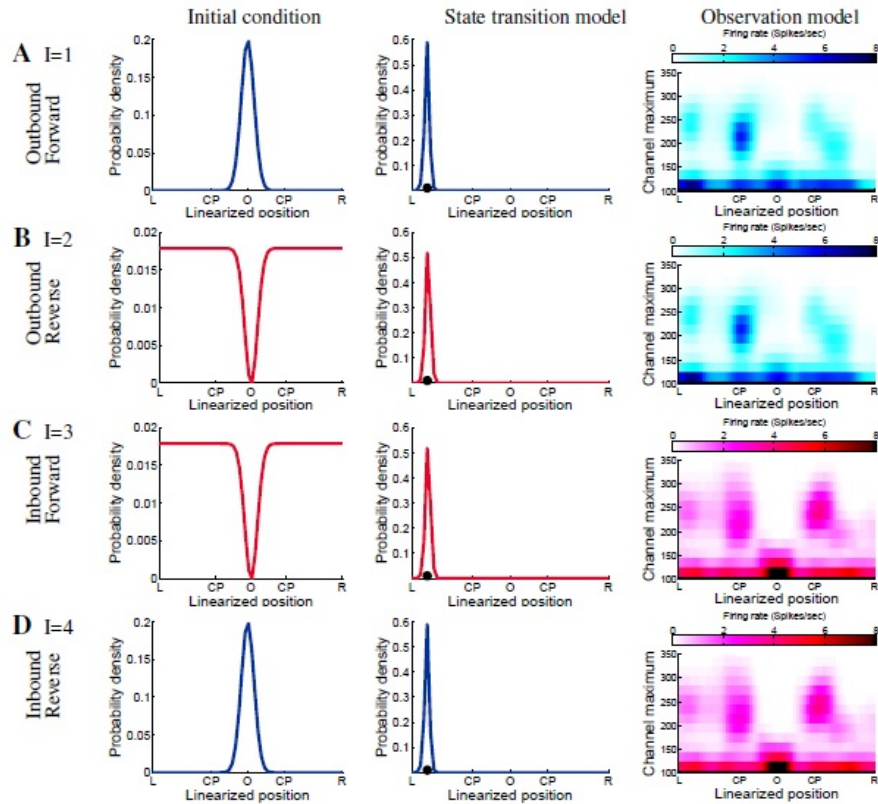


Figure 4-3: Examples of the three sources of information contributing to the decision state decoder. As a visualization aid, we plot sources of information that are consistent with those during actual outbound movement in shades of blue and sources of information that are consistent with those during actual inbound movement in shades of red. (A, B, C, D): *Left* Initial condition of replay position conditioned on the decision state. *Center* State transition model of replay position conditioned on the decision state. A slice (when the animal is at a position on the left arm shown as the black dot) of the state transition matrix is plotted here *Right* Estimated observation or likelihood model of joint mark intensity conditioned on the decision state.

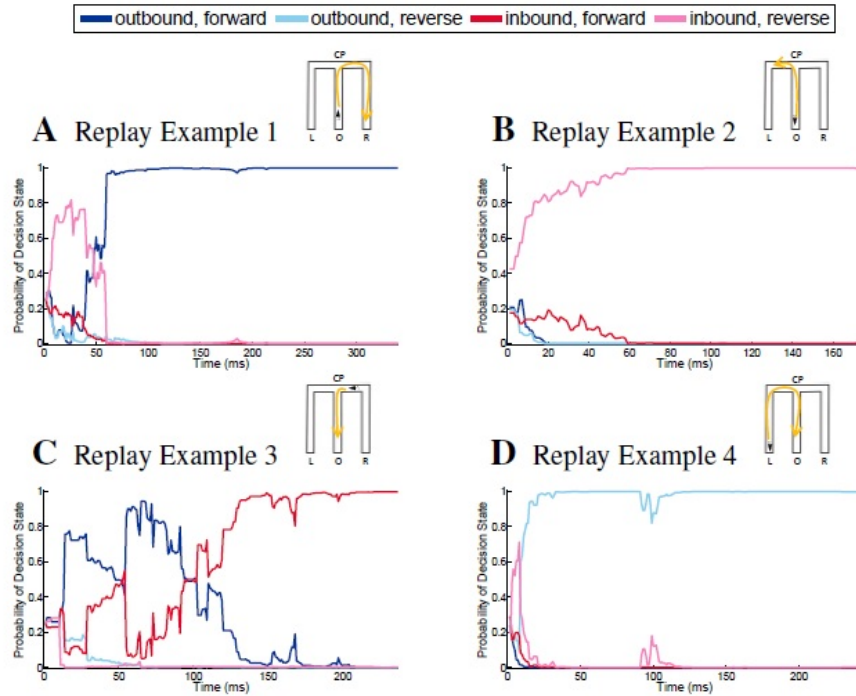


Figure 4.4: Decoding a discrete, fixed decision state. (A, B, C, D): Probability of a decision state of four example replay events as a function of time. Probability of the replay event representing an “outbound and forward” path, an “outbound and reverse” path, an “inbound and forward” path, or an “inbound and reverse” path is plotted in darker blue, lighter blue, darker red, and lighter red, respectively. *Top* Schematic representation of the decoding results where the yellow arrowed line represents the temporal evolution of the replay trajectory and the black triangle represents the replayed orientation of the rat at the beginning of the event.

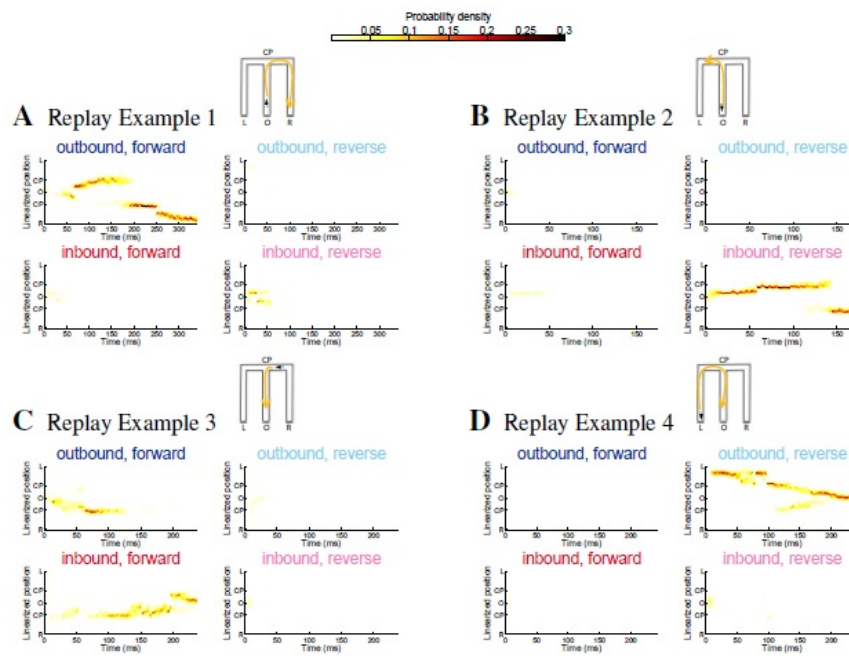


Figure 4-5: Interpreting the decision state filter as a set of weighted, parallel marked point process filters. (A, B, C, D) Decoded replay trajectory conditioning on each of the decision states for four example replay events where the heat plot shows the estimated posterior density at each time step.

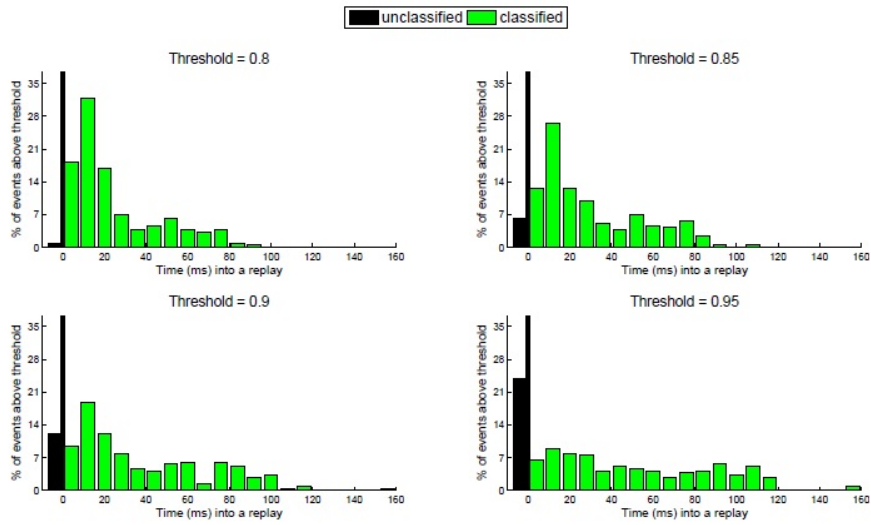


Figure 4-6: Relative frequency histogram of decision state as a function of time into a replay event, with four different thresholds on $\Pr(I)$. Events that can be classified are plotted in green. Events whose decision state probabilities did not pass the threshold and cannot be classified are plotted in black.

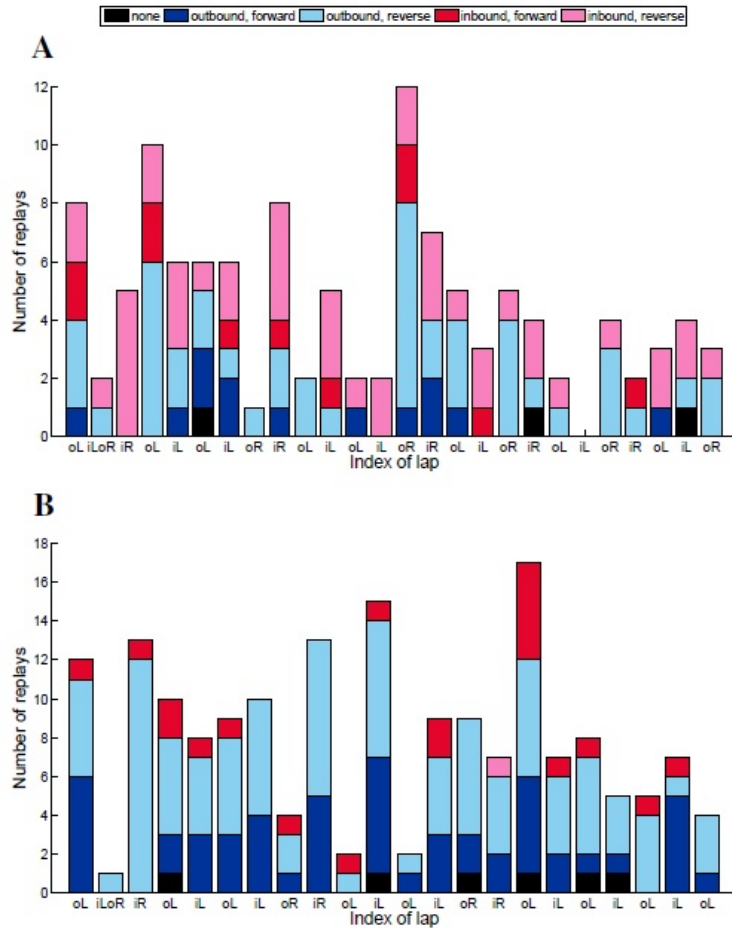


Figure 4-7: Decision state decoding results for two recording epochs on the same experiment day. Summary of decision state, by lap, with a threshold of 0.9. Replay event classified as an “outbound, forward” path, an “outbound, reverse” path, an “inbound, forward” path, or an “inbound, reverse” path is plotted in darker blue, lighter blue, darker red, and lighter red, respectively. On the axis, left and outbound lap, left and inbound lap, right and outbound lap, and right and inbound lap are denoted on the x-axis as “oL”, “iL”, “oR”, and “iR”. (A) A novel track. (B) A familiar track.

Chapter 5

Estimating a Dynamic State to Relate Neural Spiking Activity to Behavioral Signals during Cognitive Tasks

5.1 Introduction

Demonstrating the existence of meaningful relationships between behavior and neural activity is essential to our understanding of the brain and has been a subject of intensive investigation in neuroscience. Behavioral and cognitive neuroscientists aim to understand how the brain uses neural activity to integrate sensory inputs, control movements, facilitate learning and memory, activate and express emotions, etc. Neuroengineers focus on how to decode and stimulate neural activity to assist, supplement or suppress behavior.

Investigations of these questions have led to new multi-faceted experimental design and have generated behavioral data with growing complexity. In addition, technological advances now allow for recording of large quantities of information from the brain at multiple spatial and temporal scales, includes multi-channel electrode arrays, EEG, MEG and fMRI. Access to this type of high-dimensional data, both behavioral and neural, has also presented a challenge for statistical data analysis and modeling: What is an adequate representation of the relation between features of the behavioral task and structures in the neural activity?

Traditionally, studies of neural coding focus on the relation between electrophysi-

ological data and directly observable biological or behavioral signals. For example, place-field models describe spiking activity in hippocampus as a function of an animal's position in its environment. Recently, there has been increasing interest in models relating neural activity to more general variables that influence multiple aspects of behavior and cognitive function. Understanding the structure of such cognitive variables may be essential to the study of multiple neural disease processes. For example, deficits in cognitive flexibility have been linked to autism, obsessive-compulsive disorders and schizophrenia. However, cognitive flexibility is only observable through its influence on behavior, and therefore difficult to link to neural activity directly. Other examples of cognitive states that may provide meaningful links between behavior and neural activity include features like attention, affective response tendency, and approach-avoidance level. Finally, such cognitive variables are often dynamic, leading to changing behavioral outcome to stimuli through time. An important statistical challenge is to understand neural representations of these cognitive state processes and to estimate their dynamics through time.

One approach that has been successful in linking neural data to dynamic, unobserved signals is state-space modeling (Smith and Brown, 2003; Eden et al., 2004; Truccolo et al., 2005; Srinivasan et al., 2006; Czanner et al., 2008; Kemere et al., 2008; Wu et al., 2009; Paninski et al., 2010). For example, state-space modeling has been used to predict the movements of a rat from ensemble place-cell activity (Huang et al., 2009; Koyama et al., 2010). However, such methods require the behavioral signal to be estimated to be low-dimensional and directly observable during a first, encoding step. This would not be possible for the abstract, cognitive state processes, described above.

Here we present one possible solution: a general state-space paradigm to model and fit a low-dimensional cognitive state process that allows us to relate outcomes

of various behavioral tasks to simultaneously recorded neural activity across multiple brain areas. The paradigm consists of three steps. First, we estimate the dynamics of a cognitive state variable using previous knowledge of its influence on observed behavioral signals. Second, we construct models that use the estimated state and relevant covariates related to behavior to describe the statistical structure of neural activity. Last, we estimate the dynamic state again, this time using only the neural activity. We illustrate the proposed paradigm with an application to data recorded in the lateral prefrontal cortex (PFC) and caudate nucleus of non-human primates as they perform learning and adaptation in a rule-switching task.

5.2 Cognitive State-space Decoding Paradigm

In this section, we first construct a low-dimensional cognitive state process to relate behavior and neural activity. We then describe a general three-model paradigm to estimate the cognitive state in three steps, summarized in Figure 5.1.

5.2.1 Model framework

We model the underlying cognitive function (signal) during a task as a stochastic and dynamic process. The abstract state variable x_k evolves through time according to some dynamics:

$$x_k | x_{k-1} \sim f(x_k | x_{k-1}). \quad (5.1)$$

In most situations, if not all, we do not observe the cognitive state. In other words, x_k is “hidden”. Some examples of this unobserved abstract cognitive state are reward motivation level, susceptibility to fear, flexibility of learning, etc.

Instead, we obtain some behavioral data \vec{z}_k related to the task, parameterized by θ_z :

$$\vec{z}_k | x_k \sim g(\vec{z}_k | x_k; \theta_z). \quad (5.2)$$

\vec{z}_k can be a vector with components from distinct distributions, both discrete and continuous. For example, \vec{z}_k can include a binary component of correct vs. incorrect choices and a continuous component of reaction times.

We also record, in addition to the behavioral data, some electrophysiological activity \vec{y}_k , parameterized by θ_y :

$$\vec{y}_k | x_k \sim h(\vec{y}_k | x_k; \theta_y). \quad (5.3)$$

Similarly, \vec{y}_k can also include both discrete neural signals such as single unit spiking activity and continuous neural signals such as local field potentials.

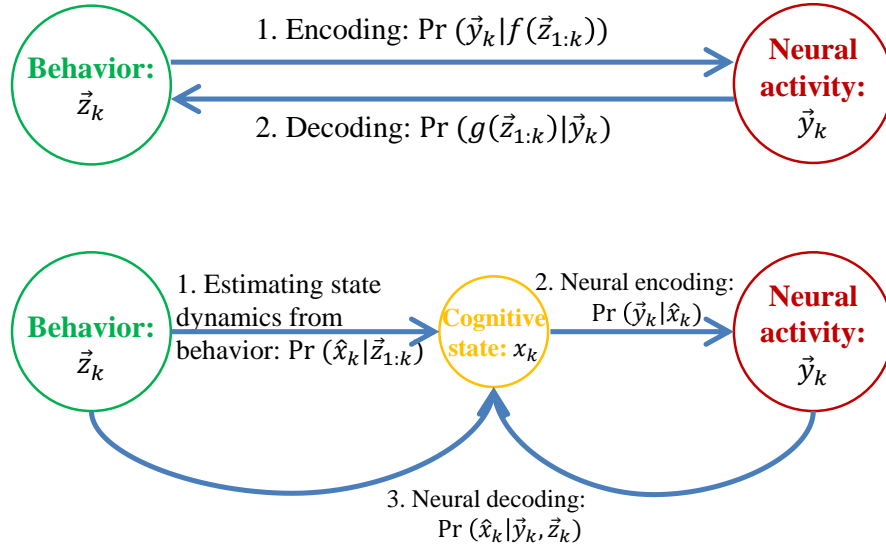


Figure 5-1: Schematic representation of the general three-model cognitive state-space paradigm (*Lower*) in comparison with the classical two-model state-space paradigm (*Upper*).

5.2.2 Estimating state dynamics from behavior

Our paradigm for estimating the cognitive state x_k involves three steps. In the first step, we estimate the state dynamics, $p(x_k|\vec{z}_{1:k})$, using previous knowledge of the behavioral task structure $g(\vec{z}_k|x_k;\theta_z)$ and some smoothing constraints $f(x_k|x_{k-1})$ on the state process dynamics. Assuming the function g in (2) is known, we use this known structure of g to estimate the state dynamics during behavioral experimental tasks:

$$p(x_k|\vec{z}_{1:k}) \propto p(\vec{z}_k|x_k) \int p(x_k|x_{k-1})p(x_{k-1}|\vec{z}_{1:k-1})dx_{k-1}. \quad (5.4)$$

The integral on the right hand side of (4) is the one-step prediction density $p(x_k|\vec{z}_{1:k-1})$ defined by the Chapman-Kolmogorov equation. Here we have assumed that given the past state value, x_{k-1} , the distribution of the current state does not depend on the past behavior. The integral in (4) typically cannot be solved analytically, but multiple numerical and approximation methods are available to compute its value at each time point. One approach is to apply some kind of filtering algorithm such as Kalman filters, which compute the distribution of the state given parameter estimates $\hat{\theta}$. Extensions of these methods, such as Expectation-Maximization (EM) and sequential Monte Carlo, simultaneously optimize the model parameter estimates $\hat{\theta}$ and the unobserved cognitive state process.

5.2.3 Neural encoding

Second, we characterize the relationship between the cognitive state and features of neural activity. In this encoding step, we use an estimate of x_k from the previous step to identify models for function $h(\vec{y}_k|x_k;\theta_y)$ with unknown structures in (3).

An example of $h(\vec{y}_k|x_k;\theta_y)$ can be a point process model with condition intensity

function $\lambda(x_k)$, when \vec{y}_t is single unit neural spiking activity:

$$p(\vec{y}_k|x_k) = [\lambda(x_k; \theta_y)\Delta]^{\vec{y}_k} \exp[\lambda(x_k; \theta_y)\Delta], \quad (5.5)$$

where $\lambda(x_k; \theta_y)$ can be estimated by parametric models of generalized linear model form.

Another example of $h(\vec{y}_k|x_k; \theta_y)$ can be multivariate Gamma models with mean $\mu(x_k; \theta_y)$, when \vec{y}_t is vector of power estimates in local field potential at specific frequencies.

In any case, we treat the estimated state process from step 1 as known, with some uncertainty, and use model fitting methods, such as maximum likelihood to estimate the unknown model parameters.

5.2.4 Neural decoding

Third, we estimate the dynamic state x_k from a new dataset that includes both neural and behavioral activity. More specifically, in this “decoding” step, we compute the posterior distribution of the state process conditioned on the observed neural activity up until the current time:

$$p(x_k|\vec{y}_k, \vec{z}_k) \propto p(\vec{y}_k, \vec{z}_k|x_k) \int p(x_k|x_{k-1})p(x_{k-1}|\vec{y}_{1:k-1}, \vec{z}_{1:k-1})dx_{k-1}. \quad (5.6)$$

If we assume that conditioning on the state, behavior and neural activity are independent, then

$$p(\vec{y}_k, \vec{z}_k|x_k) = p(\vec{y}_k|x_k)p(\vec{z}_k|x_k). \quad (5.7)$$

If we choose to decode x_k during the structured behavioral task, then we use both known g and identified models for h to estimate x_k . If we choose to decode x_k outside of the structured task, we use identified h to estimate x_k .

5.3 Application: Decoding Learning State from Spiking Activity in Monkey PFC and Caudate

5.3.1 Experimental data

We illustrate the application of the proposed paradigm with an example study. The behavioral and neural data are obtained from two monkeys performing a temporally delayed, on-line learning task in which they had to determine by trial-and-error which of four picture cues or spatial locations was currently rewarded within a learning block. Individual blocks followed either a “spatial” or an “object” rule. In the “spatial” rule, the animal was required to choose the target in the same location on every trial (e.g., always upper right). In the “object” rule, the correct action was to choose a picture that matched a picture cue (e.g., always a blue sailboat). The “spatial” rule is substantially easier and rewards perseverative behavior, while the “object” rule rewards flexibility. Details of the behavioral paradigm, data acquisition, and previous analyses of this experimental data are discussed in (Asaad and Eskandar, 2011).

In this specific example, the behavioral data is whether the monkey chose the correct location on each trial. The neural data is spiking activity recorded in the lateral prefrontal cortex (PFC) and caudate nucleus of the monkeys. The cognitive state is whether the subject has learned the rule of the task.

5.3.2 Estimate state dynamics from behavior

We take advantage of previous development of a dynamic approach to analyzing learning experiments with binary responses (Asaad and Eskandar, 2011; Smith et al., 2004; Smith et al., 2005; Suzuki and Brown, 2005; Prerau et al., 2009). We use a state-space model of learning in which a Bernoulli probability model describes behavioral task responses and a Gaussian state equation describes the hidden state process.




Figure 5.2: Two examples of the EM algorithm applied in the analysis of the state dynamics in a rule-switching task. The correct and incorrect responses are shown, respectively, by black and gray marks above the panels. The probability of a correct response occurring by chance is 0.25 in horizontal line. Solid blue lines are the learning curve estimates, and the dotted blue lines are the associated 95% confidence intervals.

In other words, $p(\vec{z}_k|x_k)$ in (4) is expressed as the Bernoulli probability mass function:

$$p(\vec{z}_k|x_k) = q_k^{\vec{z}_k} (1 - q_k)^{1 - \vec{z}_k}, \quad (5.8)$$

where q_k is defined by the logistic equation:

$$q_k = \frac{\exp(\mu + x_k)}{1 + \exp(\mu + x_k)}, \quad (5.9)$$

and μ is determined by the probability of a correct response by chance in the absence of learning or experience. Here x_k defines the learning state of the animal at trial k in the experiment. The unobservable state process $x_k|x_{k-1}$ is defined as a random walk:

$$x_k = x_{k-1} + \epsilon_k, \quad (5.10)$$

where the ϵ_k are independent Gaussian random variables with mean 0 and variance σ_ϵ^2 . The one-step prediction density $p(x_k|\vec{z}_{1:k-1})$, or learning curve, is the probability of a correct response as a function of the state process and is calculated using an EM

algorithm:

$$\begin{aligned}
 & f(q|\mu, x_{k|k}, \sigma_{k|k}^2) \\
 &= [(2\pi\sigma_{k|k}^2)^{1/2}q(1-q)]^{-1} \exp\left(-\frac{1}{2\sigma_{k|k}^2}[\log[q[(1-q)\exp(\mu)]^{-1}] - x_{k|k}]^2\right). \quad (5.11)
 \end{aligned}$$

Detailed estimation methods are referred to in (Smith et al., 2004).

Figure 5.2 shows two examples of the learning curves estimated by the EM algorithm in two learning blocks in the rule-switching behavioral task. The correct and incorrect responses are shown, respectively, by black and gray marks above the panels. Neglecting the possibility of behavioral preferences or other biases, the probability of a response occurring by chance is shown as a horizontal line at 0.25. Solid blue lines are the learning curve estimates, and the dotted blue lines are the associated 95% confidence intervals. The lower confidence bounds for the learning trial estimates remained above 0.25 after trial 14 and 63, which are, respectively the learning trials for the two learning block examples shown here.

For simplicity, we further dichotomize the trials within each learning block to be “learned” trials if the lower bound of the learning state estimate remains above 0.25 for the remainder of the trial block or “not learned” trials if otherwise.

5.3.3 Neural encoding

Because the neural data in our example is spiking activity, we present a point process generalized linear model (GLM) approach (Brillinger, 1992; Truccolo et al., 2005) for constructing a conditional intensity function that characterizes the spiking activity of PFC and caudate neurons. The conditional intensity function (Cox and Isham, 1980; Daley and Vere-Jones, 2003) relates spiking probability simultaneously to the temporal features of the behavioral task.

In this case, the conditional intensity model is defined as follows:

$$\log \lambda^c(t) = \sum_{j=1}^2 \sum_{i=1}^N \alpha_{i,j}^c B_{i,j}^c(t). \quad (5.12)$$

Here $c = 1, \dots, C$ is the index of the neuron. $j = 1, 2$ is the binary indicator of the behavioral outcome of the trial, where $j = 1$ and $j = 2$ are “learned” and “not learned” states, respectively. $B_{i,j}^c(t)$ is a basis function for a cardinal spline for neuron c , trial type j . Cardinal splines are locally defined third-order polynomial functions that flexibly approximate arbitrary smooth functions using a small number of basis functions (Ramsay and Silverman, 2010). Here, we use spline functions to capture the firing probability as a function of time relative to the picture cue. N is the number of spline control points used to fit the data. Here we chose $N = 16$ control points. $\theta = [\{\alpha_{i,j}^c\}_{i=1}^N]$ is a set of unknown parameters which relate the temporal features of the behavioral task to instantaneous spike rate.

It follows from the definition of the conditional intensity function that the probability of a spike from neuron c in a small time interval $[t, t + \Delta)$ is approximately:

$$\Pr(\text{Spike from neuron } c \text{ in } [t, t + \Delta) | \theta) \approx \lambda^c(t | \theta) \Delta. \quad (5.13)$$

This spiking intensity function describes a GLM for the spike train data. Such GLMs have a number of nice properties, including convexity of the likelihood surface and asymptotic normality of the parameter estimates, which allow us to compute maximum likelihood estimates for the model parameters in a straightforward manner. We fit these GLMs using the estimated learning state from the behavioral data. We examine the model fits to the data from 500 ms before picture cue to 2500 ms after picture cue.

Figure 5.3 shows the model parameters and their uncertainty for the maximum likelihood fit to four example neurons in this spiking data. Each panel shows the spline

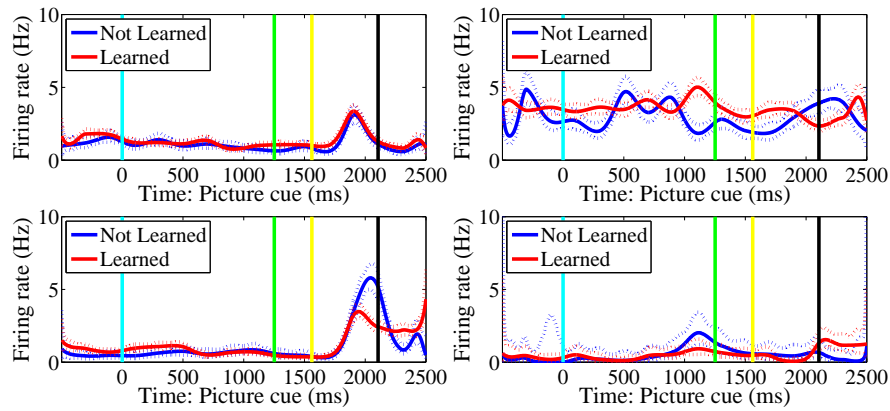


Figure 5.3: Four examples of the estimated conditional intensity with 95% confidence bounds as a function of the temporal features in the task. In each panel, blue and red solid line is the estimated conditional intensity for the “not learned” and “learned” trials, respectively. The vertical lines in cyan, green, yellow and black are the time for picture cue, go cue, feedback, and start of the inter-trial interval, respectively.

estimates, in solid lines, and 95% confidence bounds in dashed lines, as a function of time relative to picture cue, represented by the vertical line in cyan. The times of go cue, feedback, and start of inter-trial interval are identified as vertical lines in green, yellow and black, respectively. The estimated intensity and 95% confidence bounds for the learned state and not-learned state are plotted in red and blue, respectively.

Top two and the lower-left panels in Figure 5.3 show the model fit for three neurons in the caudate nucleus. For the neuron plotted in the top-left, at around 1000 ms after the picture cue and right before the go cue, the estimated intensity for learned trials in red is significantly higher than the estimated intensity for not-learned trials in blue. For the neuron plotted in the top-right, at around 2000 ms after the picture cue and within a 300 ms lag of the feedback, the estimated intensity for not-learned trials is significantly higher than that of the learned trials. For the neuron plotted in the lower-left, at around 800 ms after the picture cue, the estimated intensity for not-learned trials is significantly higher than that of the learned trials.

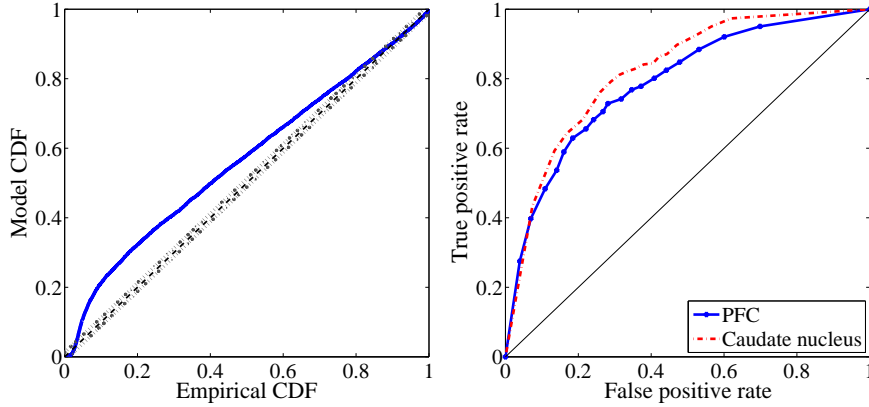


Figure 5.4: *Left:* An example Kolmogorov-Smirnov (KS) plot of time-rescaled inter-spike intervals for the estimated point process neural encoding model. It plots the empirical cumulative distribution of the rescaled ISIs against the theoretical cumulative distribution of the Exponential(1) distribution. *Right:* An example receiver-operating characteristic (ROC) curve. It plots sensitivity of the cut-off, the probability of rejecting the null hypothesis when it is false, versus significance level, the probability of rejecting a null hypothesis when it is true. ROC curves using decoding results from the neural ensemble in PFC and caudate are plotted in solid blue lines and dashed red line, respectively.

The lower-right panel in Figure 5.3 shows the model fit for a neuron in the PFC. The 95% confidence bounds for the learned and not-learned trials are always overlapping, which means that the temporal spiking properties during the observation interval are not significantly different for the two learning states.

To assess the goodness-of-fit of the two-state model, we constructed Kolmogorov-Smirnov (KS) plots of time-rescaled inter-spike intervals (ISIs) (Brown et al., 2001a). The time-rescaling theorem produces a set of rescaled ISIs that are independent with an exponential distribution with mean 1 if the proposed model accurately describes the structure in the observed spiking activity. To construct the KS plot, we plot the empirical cumulative distribution of the rescaled ISIs against the theoretical cumulative distribution of the Exponential(1) distribution. The better the quality of the model fit, the closer the K-S plot should be to a 45 degree line (Johnson and Kotz,

1980). The left panel in Figure 5-4 shows an example K-S plot for the model fit. The model and empirical CDFs demonstrate a good overall fit, with some evidence of misfit in the smaller rescaled ISIs. This may suggest some model misfit related to our assumptions of a dichotomized cognitive state and a lack of spiking history dependence structure in the point process model. More accurate modeling would likely lead to improvements in the overall goodness-of-fit and the resulting decoded state estimates.

5.3.4 Neural decoding

The previous subsection focused on the construction of neural spiking models, which uses relevant covariates related to behavior to describe the statistical structure of neural spiking activity. In this subsection we present a simple recursive Bayesian algorithm to decode the dynamic state from the spiking activity. For each trial k , we compute the posterior distribution of the monkey’s learning state given the combined spiking activity of the neural ensemble within the observation interval, $[0, T]$:

$$p(\vec{x}_k | \Delta N_{1:T}) \propto p(\Delta N_{1:T} | \vec{x}_k) p(\vec{x}_k). \quad (5.14)$$

Here $p(\vec{x}_k)$ is the prior distribution of the state. In this case, we choose a uniform prior for the binary state. The observation model, or likelihood, is given by

$$p(\Delta N_{1:T} | \vec{x}_k) \propto \prod_{t=1}^T \prod_{c=1}^C [\lambda^c(\vec{x}_k) \Delta]^{\Delta N_t} \exp(-\lambda^c(\vec{x}_k) \Delta). \quad (5.15)$$

We perform a classification procedure by thresholding the posterior of the state at various cut-off probabilities to determine whether the monkey is in a “learned” or a “not learned” state during a particular trial. The right panel in Figure 5-4 shows the receiver-operating characteristic (ROC) curve (Green and Swets, 1966; Zweig and Campbell, 1993). It plots the sensitivity of the cut-off, the probability of rejecting the

null hypothesis when it is false, versus significance level, the probability of rejecting a null hypothesis when it is true. ROC curves using decoding results from the neural ensemble in PFC and caudate are plotted as a solid blue line and a dashed red line, respectively. The ROC curve based on spiking activity in the caudate is consistently above the ROC curve based on spiking activity in PFC. It shows that neural activity in the caudate provides a better decoding of the learning state than PFC, which corroborates previous findings in the literature that the caudate contributes more closely to learning (Packard and McGaugh, 1996; Winocur and Eskes, 1998; Haruno et al., 2004; Seger and Cincotta, 2005).

5.4 Discussion

The classical two-model state-space paradigm has been successfully applied to relate behavior and neural activity directly in low-dimensional, directly observable data. However, when both the behavioral and neural data become high-dimensional and multi-faceted, this direct approach becomes computationally challenging. Here we proposed a new three-model paradigm to characterize the relationship between behavior and the neural activity. We first introduced a cognitive state process whose dynamics can be estimated from behavior. We then used the state and relevant covariates related to behavior to describe the neural activity. Lastly, we estimated the dynamic state from a combination of behavioral and neural data.

We illustrated our paradigm with a specific example of two monkeys performing a temporally delayed, on-line learning task. We demonstrated that accurate decoding of the learning state is possible with a simple point process model of population spiking. Our analyses also allowed us to compare decoding accuracy across neural population in the PFC and caudate nucleus.

Immediate extensions to the application of the paradigm shown here are under

active development. First, more accurate statistical descriptions of the behavioral data hopefully will lead to a more accurately estimated learning curve. Second, instead of working with the simplified, dichotomized learning state process, the neural encoding and decoding steps can deal directly with a continuous state process. Last, to improve the quality of fit, the point process models used for neural encoding can be expanded to include spiking history.

The essential goal of the proposed paradigm is to demonstrate the existence of meaningful relationships between complex behavior and high-dimensional neural activity. We achieve dimensionality-reduction by using hidden cognitive state processes to represent the relationship. In principle, identifying low-dimensional states that provide meaningful links between behavior and neural data is a major challenge that is specific to particular neural processing tasks. In many cases, neuroscientists may have prior conceptions of cognitive features that could represent such meaningful links. Future work may allow us to develop principled methodologies to identify cognitive states directly from data. Furthermore, by assigning some cognitive meaning to the hidden state, we can design experiments to determine the effect of manipulations of neural activity on cognitive influences of behavior. For the specific example shown here, the cognitive state can be thought of as a measure of learning flexibility, and we can modulate it to facilitate learning in the monkeys.

We envision the proposed paradigm to play a future role in the development of new types of closed-loop experiments, aiming to characterize causal relationships between neural activity and the behavior they encode. The proposed algorithm can allow investigators to identify and manipulate a low-dimensional correlate of cognitive influence in a content-specific way, altering neural activity related to certain cognitive features to modulate behavior. This may be an important step in treating mental diseases such as post-traumatic stress disorder and obsessive-compulsive disorders

clinically.

Chapter 6

Conclusions

My research addresses the relationship between the dynamics of a physical system and the occurrences of point process events related to that physical system. My graduate work focuses on problems at the intersection of statistics and neuroscience: in particular, I have proposed several new methodological developments that advance the classical framework of state-space models with point process observations, including new methods to characterize 1) the rhythmic spiking activity by including history-dependent structure, 2) to incorporate spike waveform by accommodating marked data, 3) to allow for real-time decision making, and 4) to take into account the need for dimensionality reduction for high-dimensional state and observation processes. Below, I summarize my graduate research and outline future research directions that remain to be pursued.

6.1 Contributions of the Thesis

Optimizing placement of deep brain stimulation electrodes: Full therapeutic benefit from deep brain stimulation (DBS) for Parkinson's patients critically depends on electrode location. Improving DBS targeting efficacy requires an algorithm that can reliably characterize neurophysiologic features from spiking activity. I proposed a point process modeling framework to characterize the rhythmic spiking dynamics in spike trains, test for statistically significant changes to those dynamics, and track the temporal evolution of such changes. I constructed a point process model with two sets

of parameters defining the influence of past history in two distinct physiological states. I estimated these parameters as well as the state transition times that maximized the likelihood of the observed data. I then developed a maximum likelihood ratio test to determine whether such a model provides a significant improvement over one that has only a single set of history parameters and no dynamical transitions. In case there existed a statistically significant difference between the two models, I applied a point process filtering and smoothing algorithm to track the temporal dynamics of the transition through time. My analyses uncovered a consistent dynamic pattern of modulation that characterized the structure of, and suggested potential mechanisms for, beta frequency oscillations in the spiking data.

Selectively interrupt sharp-wave ripple replay events based on replay content in real-time: Sharp-wave ripple (SWR) events in the hippocampus replay millisecond-timescale patterns of place cell activity related to the past experience of an animal. A deeper understanding of how this replayed information contributes to learning and decision-making requires the ability to manipulate SWR events based on their content. One goal is to develop a decoding algorithm to determine if a replay event's content represents a specific sequence, which will then make it possible to interrupt events based on the spatial trajectory they represent. I first developed an efficient decoding algorithm that did not require multiunit spiking waveforms to be sorted into single units by generalizing the observation model to marked point processes where the mark represents a collection of vectors that characterize information about the observed spike. In the case of tetrode recordings, for example, the mark could be a length four vector of the maximum amplitudes of the spike waveforms on each of the four electrodes at every spike time. Then I extended the state variable from a single continuous state to a joint vector that also included a discrete, fixed state variable. I made use of the underlying state process in a content-specific way

by selectively including different sources of information. This algorithm is suitable for a real-time implementation with short latencies to incorporate into content-based feedback experiments.

Developing closed-loop therapies for psychological diseases with real-time neural stimulation: An important question in neuroscience is understanding the relationship between high-dimensional electrophysiological data and complex, dynamic behavioral data. I proposed a general three-step paradigm that allows us to relate behavioral outcomes of various tasks to simultaneously recorded neural activity across multiple brain areas. I developed methods to estimate the dynamics of a cognitive state variable using previous knowledge of its influence on observed behavioral signals. Then, I constructed models that used the estimated state and relevant covariates related to behavior to describe the statistical structure of neural activity. Finally, I demonstrated that the dynamic state could be accurately decoded using only the neural activity.

6.2 Future Research

The methods developed so far can be easily extended to estimate dynamics from point process observations in other systems, such as seizure onset and progression in epileptic patients or echo-locating mechanisms in bats.

In the long term, there are three major challenges in the analysis of dynamic neural spiking data that remain to be pursued.

First, as multi-electrode recordings of neural spiking activity become increasingly available in neurophysiological studies, an essential problem is to incorporate the activity of other simultaneously active neurons in the same or in different brain regions and investigate their causal relations. What would be a good measure of such functional connectivity? What would be the estimated uncertainty associated with such a

measure? Such questions drive statisticians towards developing a framework to **track the temporal dynamics of the estimated spatial topology of a network of neurons.**

Second, having studied spiking activity from neurons in multiple brain areas, trial-to-trial variability intrigues me: The human subthalamic nucleus cells exhibit little discernible trial-to-trial variation, the place-cells in rat hippocampus have moderate variability, while the non-human primate caudate and prefrontal cortex cells are highly variable. **A better understanding of trial-to-trial variability** is not only important for its physiological significance but also for its effects on statistical procedures.

Finally, with the emerging technology such as 3D real-time microscopy of Calcium signals, **establishing a connection between Calcium spikes and action potential spikes** will allow us to extend the insights we have gained from studying action potential spikes to understand the dynamics of the brain on a much larger spatial scale in a less invasive fashion.

References

- Ackermann, M. et al. (2014). Fermi-LAT observations of the gamma-ray burst GRB 130427A. *Science*, 343:42–47.
- Aki, K. and Richards, P. G. (1980). *Quantitative Seismology*. Freeman, San Francisco, CA.
- Ambrose, E., Ambrose, B. E., and Foster, D. J. (2015). Rate of reverse, but not forward hippocampal replay increases with a relative increase in reward. In *Program No. 631.04. 2015 Neuroscience Meeting Planner*. Chicago, IL: Society for Neuroscience.
- Amirnovin, R., Williams, Z. M., Cosgrove, G. R., and Eskandar, E. N. (2004). Visually guided movements suppress subthalamic oscillations in parkinson’s disease patients. *Journal of Neuroscience*, 24:11302–11306.
- Asaad, W. F. and Eskandar, E. N. (2011). Encoding of both positive and negative reward prediction errors by neurons of the primate lateral prefrontal cortex and caudate nucleus. *Journal of Neuroscience*, 31:17772–17778.
- Ba, D., Temereanca, S., and Brown, E. N. (2014). Algorithms for the analysis of ensemble neural spiking activity using simultaneous-event multivariate point-process models. *Frontiers in Computational Neuroscience*, 8:1–13.
- Babu, G. J. and Feigelson, E. D., editors (1992). *Statistical Challenges in Modern Astronomy*. Springer, New York, NY.
- Bair, W., Koch, C., Newsome, W., and Britten, K. (1994). Power spectrum analysis of bursting cells in area MT in the behaving monkey. *Journal of Neuroscience*, 14:2870–2892.
- Bar-Gad, I., Ritov, Y., and Bergman, H. (2001). The neuronal refractory period causes a short-term peak in the autocorrelation function. *Journal of Neuroscience Methods*, 104:155–163.
- Berendse, H. W., Verbunt, J. P., Scheltens, P., van Dijk, B. W., and Jonkman, E. J. (2000). Magnetoencephalographic analysis of cortical activity in Alzheimer’s disease: a pilot study. *Clinical Neurophysiology*, 111:604–612.
- Bishop, C. (2006). *Pattern Recognition and Machine Learning*. Springer, New York.

- Bonnell, J. T. and Klebesadel, R. W. (1996). A brief history of the discovery of cosmic gamma-ray bursts. In *Gamma-ray Bursts: 3rd Huntsville Symposium*, AIP Conference Proceedings, Huntsville, AL. American Institute of Physics.
- Brillinger, D. R. (1975). The identification of point processes systems. *Annals of Probability*, 3:909–924.
- Brillinger, D. R. (1988a). Maximum likelihood analysis of spike trains of interacting nerve cells. *Biological Cybernetics*, 59:189–200.
- Brillinger, D. R. (1988b). Some statistical methods for random process data from seismology and neurophysiology. *Annals of Statistics*, 16:1–54.
- Brillinger, D. R. (1992). Nerve cell spike train data analysis: a progression of technique. *Journal of the American Statistical Association*, 87:260–271.
- Brillinger, D. R., Bryant, H. L., and Segundo, J. P. (1976). Identification of synaptic interactions. *Biological Cybernetics*, 22:213–228.
- Brillinger, D. R. and Villa, A. E. P. (1997). Assessing connections in networks of biological neurons. In *The Practice of Data Analysis: Essays in Honor of John W. Tukey*, pages 77–92. Princeton University Press, Princeton, NJ.
- Brown, E. N. (2007). Signal processing and statistical challenges in neuroscience data. *SIAM News*, 40:13–14.
- Brown, E. N., Barbieri, R., Ventura, V., and Kass, R. E. (2001a). The time-rescaling theorem and its application to neural spike train data analysis. *Neural Computation*, 14:325–346.
- Brown, E. N., Barbieri, R., Ventura, V., Kass, R. E., and Frank, L. M. (2002). The time-rescaling theorem and its application to neural spike train data analysis. *Neural Computation*, 14:325–346.
- Brown, E. N., Frank, L. M., Tang, D., Quirk, M. C., and Wilson, M. A. (1998). A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *Journal of Neuroscience*, 18:7411–25.
- Brown, E. N., Mitra, P. P., and Kass, R. E. (2004). Multiple neural spike train data analysis: State-of-the-art and future challenges. *Nature Neuroscience*, 7:456–461.
- Brown, P., Oliviero, A., Mazzone, P., Insola, A., Tonali, P., and Di Lazzaro, V. (2001b). Dopamine dependency of oscillations between subthalamic nucleus and pallidum in Parkinson’s disease. *Journal of Neuroscience*, 21:1033–1038.

- Buzsaki, G. (1989). Tow-stage model of memory trace formation: A role for “noisy” brain states. *Neuroscience*, 31:551–570.
- Buzsaki, G. (2015). Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus*, 25:1073–1188.
- Carr, M. F., Jadhav, S. P., and Frank, L. M. (2011). Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nature Neuroscience*, 14:147–153.
- Casella, G. and Berger, R. L. (2001). *Statistical Inference*. Duxbury Press, Pacific Grove, CA.
- Chen, Z., Kloosterman, F., Layton, S., and Wilson, M. A. (2012). Transductive neural decoding for unsorted neuronal spikes of rat hippocampus. In *Proceedings of the 34th Annual International Conference of the IEEE EMBS*, San Diego, CA.
- Cheng, S. and Frank, L. M. (2008). New experiences enhance coordinated neural activity in the hippocampus. *Neuron*, 57:303–313.
- Cho, R. Y., Konecky, R. O., and Carter, C. S. (2006). Impairments in frontal cortical gamma synchrony and cognitive control in schizophrenia. *Proceedings of the National Academy of Sciences USA*, 103:19878–19883.
- Coben, L. A., Danziger, W. L., and Berg, L. (1983). Frequency analysis of the resting awake EEG in mild senile dementia of Alzheimer type. *Electroencephalography and Clinical Neurophysiology*, 55:372–380.
- Cox, D. R. (1955). Some statistical methods connected with series of events. *Journal of the Royal Statistical Society: Series B*, 17:129–164.
- Cox, D. R. and Isham, V. (1980). *Point Processes*. Chapman and Hall, London, UK.
- Csicsvari, J., Hirase, H., Czurko, A., Mamiya, A., and Buzsaki, G. (1999). Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving rat. *Journal of Neuroscience*, 19:274–287.
- Csicsvari, J., O’Neill, J., Allen, K., and Senior, J. (2007). Place-selective firing contributes to the reverse-order reactivation of CA1 pyramidal cells during sharp waves in open-field exploration. *European Journal of Neuroscience*, 26:704–716.
- Czanner, G., Eden, U. T., Wirth, S., Yanike, M., Suzuki, W. A., and Brown, E. N. (2008). Analysis of between-trial and within-trial neural spiking dynamics. *Journal of Neurophysiology*, 99:2672–2693.

- Daley, D. and Vere-Jones, D. (2003). *An Introduction to the Theory of Point Processes*. Springer-Verlag, New York.
- Deng, X., Liu, D. F., Kay, K., Frank, L. M., and Eden, U. T. (2015). Clusterless decoding of position from multiunit activity using a marked point process filter. *Neural Computation*, 27:1438–1460.
- Diba, K. and Buzsaki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, 10:1241–1242.
- Dobigeon, N., Tourneret, J.-Y., and Scargle, J. D. (2007). Joint segmentation of multivariate astronomical time series: Bayesian sampling with a hierarchical model. *IEEE Transactions on Signal Processing*, 55:414–423.
- Donoghue, J. (2002). Connecting cortex to machines: Recent advances in brain interfaces. *Nature Neuroscience*, 5:1085–1088.
- Doucet, A., Freitas, N., and Gordon, N. (2001). *Sequential Monte Carlo Methods in Practice*. Springer, New York.
- Eden, U. T. and Brown, E. N. (2008). Continuous-time filters for state estimation from point process models of neural data. *Statistica Sinica*, 18:1293–1310.
- Eden, U. T., Frank, L. M., Barbieri, R., Solo, V., and N., B. E. (2004). Dynamic analysis of neural encoding by point process adaptive filtering. *Neural Computation*, 16:971–998.
- Eden, U. T., Gale, J. T., Amirnovin, R., and Eskandar, E. N. (2012). Characterizing the spiking dynamics of subthalamic nucleus neurons in parkinson’s disease using generalized linear models. *Frontiers in Integrative Neuroscience*, 6.
- Ego-Stengel, V. and Wilson, M. A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus*, 20:1–10.
- Ergun, A., Barbieri, R., Eden, U. T., Wilson, M. A., and Brown, E. N. (2007). Construction of point process adaptive filter algorithms for neural systems using sequential monte carlo. *IEEE Transactions on Biomedical Engineering*, 54:419–28.
- Foffani, G., Ardolino, G., Meda, B., Egidi, M., Rampini, P., Caputo, E., Baselli, G., and Priori, A. (2005). Altered subthalamo-pallidal synchronisation in Parkinsonian dyskinesias. *Journal of Neurology, Neurosurgery and Psychiatry*, 76:426–428.
- Foster, D. J. and Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440:680–683.

- Frank, L., Eden, U., Solo, V., Wilson, M., and Brown, E. (2002). Contrasting patterns of receptive field plasticity in the hippocampus and the entorhinal cortex: An adaptive filtering approach. *Journal of Neuroscience*, 22:3817–3830.
- Frank, L. M., Brown, E. N., and Wilson, M. A. (2001). A comparison of the firing properties of putative excitatory and inhibitory neurons from CA1 and the entorhinal cortex of the awake behaving rat. *Journal of Neurophysiology*, 86:2029–2040.
- Franklin, J. and Bair, W. (1995). The effect of a refractory period on the power spectrum of neuronal discharge. *SIAM Journal of Applied Mathematics*, 55:1074–1093.
- Fraser, G. W., Chase, S. M., A., W., and B., S. A. (2009). Control of a brain-computer interface without spike sorting. *Journal of Neural Engineering*, 6:055004.
- Gerhard, F., Kispersky, T., Gutierrez, G. J., Marder, E., Kramer, M. A., and Eden, U. T. (2013). Successful reconstruction of a physiological circuit with known connectivity from spiking activity alone. *PLoS Computational Biology*, 9.
- Girardeau, G., Benchenane, K., Wiener, S. I., Buzsaki, G., and Zugaro, M. B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, 12:1222–1223.
- Green, D. M. and Swets, J. A. (1966). *Signal Detection Theory and Psychophysics*. John Wiley and Sons, New York, NY.
- Grosenick, L., Marshel, J. H., and Deisseroth, K. (2015). Closed-loop and activity-guided optogenetic control. *Neuron*, 86:106–139.
- Haenschel, C., Bittner, R. A., Waltz, J., Haertling, F., Wibral, M., Singer, W., Linden, D. E., and Rodriguez, E. (2009). Cortical oscillatory activity is critical for working memory as revealed by deficits in early-onset schizophrenia. *Journal of Neuroscience*, 29:9481–9489.
- Harris, K. D., Henze, D. A., Csicsvari, J., Hirase, H., and Buzsaki, G. (2000). Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *Journal of Neurophysiology*, 84:401–414.
- Haruno, M., Kuroda, T., Doya, K., Toyama, K., Kimura, M., Samejima, K., Imamizu, H., and Kawato, M. (2004). A neural correlate of reward-based behavioral learning in caudate nucleus: a functional magnetic resonance imaging study of a stochastic decision task. *Journal of Neuroscience*, 24:1660–1665.
- Hasselmo, M. E. (2009). A model of episodic memory: mental time travel along encoded trajectories using grid cells. *Neurobiology of Learning and Memory*, 92:559–573.

- Hawkes, A. G. (1971). Point spectra of some mutually exciting point processes. *Journal of the Royal Statistical Society: Series B*, 33:438–443.
- Haykin, S. (1996). *Adaptive filter theory*. Prentice Hall, Englewood Cliffs, NJ.
- Henshaw, J. (1998). Batse gamma-ray burst, trigger 551. <http://www.synapse9.com/batse551.htm>, accessed on 2016-03-24.
- Hill, D. P., Reasenber, P. A., Michael, A., Arabaz, W. J., Beroza, G., Brumbaugh, D., Brune, J. N., Castro, R., Davis, S., dePolo, D., Ellsworth, W. L., Gomberg, J., Harmsen, S., House, L., Jackson, S. M., Johnston, M. J. S., Jones, L., Keller, R., Malone, S., Munguia, L., Nava, S., Pechmann, J. C., Sanford, A., Simpson, R. W., Smith, R. B., Stark, M., Stickney, M., Vidal, A., Walter, S., Wong, S., and Zollweg, J. (1993). Seismicity remotely triggered by the magnitude 7.3 Landers, California, earthquake. *Science*, 260:1617–1622.
- Hough, S. (2010). *Predicting the Unpredictable: The Tumultuous Science of Earthquake Prediction*. Princeton University Press, Princeton, NJ.
- Huang, C., Wahlund, L., Dierks, T., Winblad, J. B., and Jelic, V. (2000). Discrimination of Alzheimer’s disease and mild cognitive impairment by equivalent EEG sources: a cross-sectional and longitudinal study. *Clinical Neurophysiology*, 111:1961–1967.
- Huang, Y., Brandon, M. P., Griffin, A. L., Hasselmo, M. E., and Eden, U. T. (2009). Decoding movement trajectories through a t-maze using point process filters applied to place field data from rat hippocampal region ca1. *Neural Computation*, 21:3305–3334.
- Hurtado, J. M., Rubchinsky, L. L., and Sigvardt, K. A. (2004). Statistical method for detection of phase-locking episodes in neural oscillations. *Journal of Neurophysiology*, 91:1883–1898.
- Hurtado, J. M., Rubchinsky, L. L., Sigvardt, K. A., Wheelock, V. L., and Papps, C. T. E. (2005). Temporal evolution of oscillations and synchrony in GPi/muscle pairs in Parkinson’s disease. *Journal of Neurophysiology*, 93:1569–1584.
- Jadhav, S. P., Kemere, C., German, P. W., and Frank, L. M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. *Science*, 336:1454–1458.
- Jarvis, M. R. and Mitra, P. P. (2001). Sampling properties of the spectrum and coherency of sequences of action potentials. *Neural Computation*, 13:717–749.
- Jensen, A. (1948). An elucidation of Erlang’s statistical works through the theory of stochastic processes. In *The Life and Works of A. K. Erlang*, pages 23–100.

- Ji, D. and Wilson, M. A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature Neuroscience*, 10:100–107.
- Johnson, A. and Kotz, S. (1980). *Distributions in Statistics: Continuous Univariate Distributions*. Wiley, New York, NY.
- Kagan, Y. Y. and Jackson, D. (1994). Long-term probabilistic forecasting of earthquakes. *Journal of Geophysical Research*, 99:13685–13700.
- Karlsson, M. P. and Frank, L. M. (2008). Network dynamics underlying the formation of sparse, informative representations in the hippocampus. *Journal of Neuroscience*, 28:14271–14281.
- Karlsson, M. P. and Frank, L. M. (2009). Awake replay of remote experiences in the hippocampus. *Nature Neuroscience*, 12:913–918.
- Karr, A. F. (1991). *Point Process and Their Statistical Inference*. Marcel Dekker, Inc., New York, NY.
- Kemere, C., Santhanam, G., Yu, B. M., Afshar, A., Ryu, S. I., Meng, T. H., and Shenoy, K. V. (2008). Detecting neural-state transitions using hidden Markov models for motor cortical prostheses. *Journal of Neurophysiology*, 100:2441–2452.
- Khintchine, A. Y. (1956a). On Poisson streams of events. *Theory of Probability and Its Applications*, 1:248–255.
- Khintchine, A. Y. (1956b). Streams of events without aftereffects. *Theory of Probability and Its Applications*, 1:1–15.
- Klebesadel, R. W., Strong, I. B., and Olson, R. A. (1973). Observations of gamma-ray bursts of cosmic origin. *Astrophysical Journal*, 182:L85–L88.
- Kloosterman, F., Layton, S. P., Chen, Z., and Wilson, M. A. (2014a). Bayesian decoding using unsorted spikes in the rat hippocampus. *Journal of Neurophysiology*, 111:217–227.
- Kloosterman, F., Layton, S. P., Chen, Z., and Wilson, M. A. (2014b). Bayesian decoding using unsorted spikes in the rat hippocampus.
- Kolaczyk, E. D. (1999). Bayesian multi-scale models for Poisson processes. *Journal of the American Statistical Association*, 94:920–933.
- Kolaczyk, E. D. and Dixon, D. D. (2000). Nonparametric estimation of intensity maps using Haar wavelets and poisson noise characteristics. *The Astrophysical Journal*, 534:490–505.

- Kolaczyk, E. D. and Nowak, R. D. (1999). Deconvolution of poisson-limited data using a Bayesian multi-scale model. *Bulletin of the American Astronomical Society*, 31:734.
- Kolaczyk, E. D. and Nowak, R. D. (2005). Multiscale generalised linear models for nonparametric function estimation. *Biometrika*, 92:119–133.
- Koyama, S., Eden, U. T., Brown, E. N., and Kass, R. E. (2010). Bayesian decoding of neural spike trains. *Annals of the Institute of Statistical Mathematics*, 62:37–59.
- Le Cam, L. (1960). Locally asymptotically normal families of distributions. *University of California Publications in Statistics*, 3:37–98.
- Lee, A. K. and Wilson, M. A. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron*, 36:1183–1194.
- Lepage, K. Q., Kramer, M. A., and Eden, U. T. (2011). The dependence of spike field coherence on expected intensity. *Neural Computation*, 23:2209–2241.
- Levy, R., Ashby, P., Hutchison, W. D., Lang, A. E., Lozano, A. M., and Dostrovsky, J. O. (2002). Dependence of subthalamic nucleus oscillations on movement and dopamine in Parkinson’s disease. *Brain*, 125:1196–1209.
- Lewicki, M. S. (1998). A review of methods for spike sorting: the detection and classification of neural action potentials. *Network: Computation in Neural Systems*, 9:R53–R78.
- Liptser, R. S. and Shiryaev, A. N. (1978). *Statistics of Random Processes. II. Applications*. Springer.
- Luczak, A. and Narayanan, N. S. (2005). Spectral representation—analyzing single-unit activity in extracellularly recorded neuronal data without spike sorting. *Journal of Neuroscience Methods*, 144:53–61.
- Marsden, J. F., Limousin-Dowsey, P., Ashby, P., Pollak, P., and Brown, P. (2001). Subthalamic nucleus, sensorimotor cortex and muscle interrelationships in Parkinson’s disease. *Brain*, 124:378–388.
- Maselli, A. et al. (2014). GRB 130427A: A nearby ordinary monster. *Science*, 343:48–51.
- McCarthy, M. M., Moore-Kochlacs, C., Gu, X., Boyden, E. S., Han, X., and Kopell, N. (2011). Striatal origin of the pathologic beta oscillations in Parkinson’s disease. *Proceedings of the National Academy of Sciences USA*, 108:11620–11625.

- Meegan, C. A., Fishman, G. J., Wilson, R. B., Horack, J. M., Brock, M. N., Paciasas, W. S., Pendleton, G. N., and Kouveliotou, C. (1992). Spatial distribution of gamma-ray bursts observed by BATSE. *Nature*, 355:143–145.
- Meng, L., Kramer, M. A., and Eden, U. T. (2011). A sequential Monte Carlo approach to estimate biophysical neural models from spikes. *Journal of Neural Engineering*, 8.
- Meng, L., Kramer, M. A., Middleton, S. J., Whittington, M. A., and Eden, U. T. (2014). A unified approach to linking experimental, statistical and computational analysis of spike train data. *PLoS ONE*, 9.
- Mink, J. W. (1996). The basal ganglia: Focused selection and inhibition of competing motor programs. *Progress in Neurobiology*, 50:381–425.
- Mink, J. W. and Thach, W. T. (1993). Basal ganglia intrinsic circuits and their role in behavior. *Current Opinion in Neurobiology*, 3:950–957.
- Muller, R. U. and Kubie, J. K. (1989). The firing of hippocampal place cells predicts the future position of freely moving rats. *Journal of Neuroscience*, 9:4101–4110.
- Musmeci, F. and Vere-Jones, D. (1992). A space-time clustering model for historical earthquakes. *Annals of the Institute of Statistical Mathematics*, 44:1–11.
- Ogata, Y. (1978). The asymptotic behaviour of maximum likelihood estimators for stationary point process. *Annals of the Institute of Statistical Mathematics*, 30:243–261.
- Ogata, Y. (1988). Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical Association*, 83:9–27.
- Ogata, Y. (1998). Space-time point process models for earthquake occurrences. *Annals of the Institute of Statistical Mathematics*, 50:379–402.
- Ogata, Y., Akaike, H., and Katsura, K. (1982). The application of linear intensity models to the investigation of causal relations between a point process and another point process. *Annals of the Institute of Statistical Mathematics*, 40:29–39.
- O’Keefe, J. (1979). A review of the hippocampal place cells. *Progress in Neurobiology*, 13:419–439.
- O’Keefe, J. and Dostrovsky, J. (1971). The hippocampus as a spatial map. preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 34:171–175.

- Osipova, D., Ahveninen, J., Jensen, O., Ylikoski, A., and Pekkonen, E. (2005). Altered generation of spontaneous oscillations in Alzheimer's disease. *Neuroimage*, 27:835–841.
- Packard, M. G. and McGaugh, J. L. (1996). Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiology of Learning and Memory*, 65:65–72.
- Palm, C. (1943). Intesitatsschwankungen im fernsprechverkehr. *Ericsson Technics*, 44.
- Paninski, L., Ahmadian, Y., Ferreira, D. G., Koyama, S., Rad, K. R., Vidne, M., Vogelstein, J., and Wu, W. (2010). A new look at state-space models for neural data. *Journal of Computational Neuroscience*, 29:107–126.
- Penttila, M., Partanen, J. V., Soininen, H., and Riekkinen, P. J. (1985). Quantitative analysis of occipital EEG in different stages of Alzheimer's disease. *Electroencephalography and Clinical Neurophysiology*, 60:1–6.
- Pfeiffer, B. E. and Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, 497:74–79.
- Pfurtscheller, G. (1981). Central beta rhythm during sensorimotor activities in man. *Electroencephalography and Clinical Neurophysiology*, 51:253–264.
- Prerau, M. J., Hartnack, K. E., Obregon-Henao, G., Sampson, A., Merlino, M., Gannon, K., Bianchi, M. T., Ellenbogen, J. M., and Purdon, P. L. (2014). Tracking the sleep onset process: An empirical model of behavioral and physiological dynamics. *PLoS Computational Biology*, 10.
- Prerau, M. J., Smith, A. C., Eden, U. T., Kubota, Y., Yanike, M., Suzuki, W., Graybiel, A. M., and Brown, E. N. (2009). Characterizing learning by simultaneous analysis of continuous and binary measures of performance. *Journal of Neurophysiology*, 102:3060–3072.
- Pulliam, R. J. and Stark, P. B. (1993). Bumps on the core-mantle boundary—are they facts or artifacts? *Journal of Geophysical Research: Solid Earth*, 98:1943–1955.
- Quiroga, R. Q. (2012). Spike sorting. *Current Biology*, 22:R45–R46.
- Ramsay, J. O. and Silverman, B. W. (2010). *Functional Data Analysis*. Springer, New York, NY, 2nd edition.
- Redish, A. D. (1999). *Beyond the Cognitive Map: From Place Cells to Episodic Memory*. MIT Press, Cambridge, MA.

- Redish, A. D. and Touretzky, D. S. (1998). The role of the hippocampus in solving the morris water maze. *Neural Computation*, 10:73–111.
- Rivlin-Etzion, M., Ritov, Y., Heimer, G., Bergman, H., and Bar-Gad, I. (2006). Local shuffling of spike trains boosts the accuracy of spike train spectral analysis. *Journal of Neurophysiology*, 95:3245–3256.
- Rubinstein, J. L. and Beroza, G. C. (2007). Full waveform earthquake location: Application to seismic streaks on the Calaveras Fault, California. *Journal of Geophysical Research: Solid Earth*, 112:B05303.
- Salenius, S., Avikainen, S., Kaakkola, S., Hari, R., and Brown, P. (2002). Defective cortical drive to muscle in Parkinson’s disease and its improvement with levodopa. *Brain*, 125:491–500.
- Sanders, C. O. (1993). Interaction of the San Jacinto and San Andreas fault zones, southern California: Triggered earthquake migration and coupled recurrence intervals. *Science*, 260:973–976.
- Sarma, S. V., Cheng, M. L., Williams, Z. M., Hu, R., Eskandar, E. N., and Brown, E. N. (2009). Using point process models to determine the impact of visual cues on basal ganglia activity and behavior of Parkinsons patients. In *Proceedings of Joint 48th IEEE Conference on Decision and Control and 28th Chinese Control Conference*, Shanghai, P.R. China.
- Sarma, S. V., Eden, U. T., Cheng, M. L., Williams, Z. M., Hu, R., Eskandar, E. N., and Brown, E. N. (2010). Using point process models to compare neural spiking activity in the subthalamic nucleus of Parkinson’s patients and a healthy primate. *IEEE Transactions on Biomedical Engineering*, 57:1297–1305.
- Scargle, J. (1998). Studies in astronomical time series analysis. v. bayesian blocks, a new method to analyze structure in photon counting data. *The Astrophysical Journal*, 504:405–418.
- Scargle, J. D. and Babu, G. J. (2003). Point processes in astronomy: Exciting events in the universe. In *Handbook of Statistics*, volume 21. Elsevier Science B. V., New York, NY.
- Scargle, J. D., Norris, J. P., Jackson, B., and Chiang, J. (2013). Studies in astronomical time series analysis. vi. bayesian blocks representations. *The Astrophysical Journal*, 764:167–192.
- Schaff, D. P., Bokelmann, G. H. R., Beroza, G. C., Waldhauser, F., and Ellsworth, W. L. (2002). High-resolution image of Calaveras Fault seismicity. *Journal of Geophysical Research*, 107:2186–2201.

- Schmiedt, C., Brand, A., Hildebrandt, H., and Basar-Eroglu, C. (2005). Event-related theta oscillations during working memory tasks in patients with schizophrenia and healthy controls. *Cognitive Brain Research*, 25:936–947.
- Schreiter-Gasser, U., Gasser, T., and Ziegler, P. (1993). Quantitative EEG analysis in early onset Alzheimer’s disease: a controlled study. *Electroencephalography and Clinical Neurophysiology*, 86:15–22.
- Seger, C. A. and Cincotta, C. M. (2005). The roles of the caudate nucleus in human classification learning. *Journal of Neuroscience*, 25:2941–2951.
- Siegle, J. H. and Wilson, M. A. (2014). Enhancement of encoding and retrieval functions through theta phase-specific manipulation of hippocampus. *eLife*, 3.
- Silva, D., Feng, T., and Foster, D. J. (2015). Trajectory events across hippocampal place cells require previous experience. *Nature Neuroscience*, 18:1772–1779.
- Singer, A. C. and Frank, L. M. (2009). Rewarded outcomes enhance reactivation of experience in the hippocampus. *Neuron*, 64:910–921.
- Smith, A. C. and Brown, E. N. (2003). Estimating a state-space model from point process observations. *Neural Computation*, 15:965–991.
- Smith, A. C., Frank, L. M., Wirth, S., Yanike, M., Hu, D., Kubota, Y., Graybiel, A. M., Suzuki, W. A., and Brown, E. N. (2004). Dynamic analysis of learning in behavioral experiments. *Journal of Neuroscience*, 24:447–461.
- Smith, A. C., Stefani, M. R., Moghaddam, B., and Brown, E. N. (2005). Analysis and design of behavioral experiments to characterize population learning. *Journal of Neurophysiology*, 93:1776–1792.
- Srinivasan, L., Eden, U. T., Willsky, A. S., and Brown, E. N. (2006). A state-space analysis for reconstruction of goal-directed movement using neural signals. *Neural Computation*, 18:2465–2494.
- Stark, E. and Abeles, M. (2007). Predicting movement from multiunit activity. *Journal of Neuroscience*, 27:8387–8397.
- Suzuki, W. A. and Brown, E. N. (2005). Behavioral and neurophysiological analyses of dynamic learning processes. *Behavioral and Cognitive Neuroscience Reviews*, 4:67–95.
- Todorova, S., Sadtler, P., Batista, A., Chase, S., and Ventura, V. (2014). To sort or not to sort: the impact of spike-sorting on neural decoding performance. *Journal of Neural Engineering*, 11:056005.

- Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue, J. P., and Brown, E. N. (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of Neurophysiology*, 94:1074–1089.
- Tulving, E. and Markowitsch, H. J. (1998). Episodic and declarative memory: role of the hippocampus. *Hippocampus*, 8:198–204.
- Udias, A. (1989). Development of fault-plane studies for mechanism of earthquakes. In *Observatory Seismology*. University of California Press, Berkeley, CA.
- Uhlhaas, P. J. and Singer, W. (2010). Abnormal neural oscillations and synchrony in schizophrenia. *Nature Review Neuroscience*, 11:100–113.
- van der Vaart, A. W. (1998). *Asymptotic Statistics*. Cambridge University Press, New York, NY.
- Ventura, V. (2008). Spike train decoding without spike sorting. *Neural Computation*, 20:923–963.
- Ventura, V. (2009a). Automatic spike sorting using tuning information. *Neural Computation*, 21:2466–2501.
- Ventura, V. (2009b). Traditional waveform based spike sorting yields biased rate code estimates. *Proceedings of the National Academy of Sciences USA*, 106:6921–6.
- Vere-Jones, D. (1995). Forecasting earthquakes and earthquake risk. *International Journal of Forecasting*, 11:503–538.
- Wesson, R. L., Bakun, W. H., and Perkins, D. M. (2003). Association of earthquakes and faults in the San Francisco bay area using Bayesian inference. *Bulletin of the Seismological Society of America*, 93:1306–1332.
- Wild, J., Prekopcsak, Z., Sieger, T., Novak, D., and Jech, R. (2012). Performance comparison of extracellular spike sorting algorithms for single-channel recordings. *Journal of Neuroscience Methods*, 203:369–376.
- Wilson, M. A. and McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science*, 261:1055–1058.
- Winocur, G. and Eskes, G. (1998). Prefrontal cortex and caudate nucleus in conditional associative learning: dissociated effects of selective brain lesions in rats. *Behavioral Neuroscience*, 112:89–101.

- Winterer, G., Ziller, M., Dorn, H., Frick, K., Mulert, C., Wuebben, Y., Herrmann, W. M., and Coppola, R. (2000). Schizophrenia: reduced signal-to-noise ratio and impaired phase-locking during information processing. *Clinical Neurophysiology*, 111:837–849.
- Wold, A. (1948). On stationary point processes and Markov chains. *Scandinavian Actuarial Journal*, 31:229–240.
- Wu, W., Gao, Y., Bienenstock, E., Donoghue, J. P., and Black, M. J. (2006). Bayesian population coding of motor cortical activity using a Kalman filter. *Neural Computaiton*, 18:80–118.
- Wu, W., Kulkarni, J., Hatsopoulos, N., and Paninski, L. (2009). Neural decoding of goal-directed movements using a linear statespace model with hidden states. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 17(4).
- Wu, X. and Foster, D. J. (2014). Hippocampal replay captures the unique topological structure of a novel environment. *Journal of Neuroscience*, 34:6459–6469.
- Young, C., Meredith, D. C., and Ryan, J. M. (1995). A compact representation of gamma-ray burst time series. *Astrophysics and Space Science*, 231:119–122.
- Zhang, K., Ginzburg, I., McNaughton, B. L., and Sejnowski, T. J. (1998). Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *Journal of Neurophysiology*, 79:1017–1044.
- Zweig, M. H. and Campbell, G. (1993). Receiver-operating characteristic (ROC) plots: a fundamental evaluation tool in clinical medicine. *Clinical Chemistry*, 39:561–577.

Curriculum Vitae

- Contact* Xinyi Deng
 Department of Mathematics and Statistics, Boston University, 111
 Cummington Mall, Boston, MA 02215, USA
- Education* **American University**, B.S., Statistics, *Magna Cum Laude*, Au-
 gust 2007–May 2011
Boston University, Ph.D. candidate, August 2011–present. The-
 sis advisor: Uri T. Eden.
- Publications*
1. Deng X, Faghih RT, Barbieri R, Paulk AC, Asaad WF, Brown EN, Dougherty DD, Widge AS, Eskandar EN, and Eden UT. (2015) Estimating a dynamic state to relate neural spiking activity to behavioral signals during cognitive tasks. *Proceedings of the 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society* (pp. 7808–7813). Milan, Italy.
 2. Deng X, Liu DF, Kay K, Frank LM, and Eden UT. (2015) Clusterless decoding of position from multiunit activity using a marked point process filter. *Neural Computation*, 27(7), 1438–1460.
 3. Deng X, Eskandar EN, and Eden UT. (2013) A point process approach to identifying and tracking transitions in neural spiking dynamics in the subthalamic nucleus of Parkinson’s patients. *Chaos*, 23(4), 046102.
 4. Prado-Oviedo NA, Malloy EJ, Deng X, and Brown JL. (2013) Hyperprolactinemia is not associated with hyperestrogenism in non-cycling African elephants (*Loxodonta africana*). *General and Comparative Endocrinology*, 189, 7–14.