

2014

Advances in point process filters and their application to sympathetic neural activity

<https://hdl.handle.net/2144/15226>

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

BOSTON UNIVERSITY
GRADUATE SCHOOL OF ARTS AND SCIENCES

Dissertation

**ADVANCES IN POINT PROCESS FILTERS
AND
THEIR APPLICATION TO SYMPATHETIC NEURAL ACTIVITY**

by

YEVGENIY ZAYDENS

B.S., Rutgers University, 2000
M.A., Boston University, 2008

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

2014

© 2014 by
YEVGENIY ZAYDENS
All rights reserved except for Chapter 2,
which is © 2013 The Institute of Electrical
and Electronics Engineers, Inc.

Approved by

First Reader

Uri T. Eden, Ph.D.
Associate Professor of Mathematics and Statistics
Boston University

Second Reader

J. Andrew Taylor, Ph.D.
Professor of Physical Medicine and Rehabilitation
Harvard University, School of Medicine

Third Reader

Michael A. Cohen, Ph.D.
Associate Professor of Computer Science
Boston University

ACKNOWLEDGMENTS

First and foremost, I would like to express my gratitude to my primary advisor, Prof. Uri Eden. His extensive knowledge and intuition inspired me to reach higher than I thought possible, and his patient guidance allowed me to get there. Uri was a true mentor and a friend, someone I could rely on both professionally and personally.

I would also like to thank my committee for their invaluable advice, scientific resourcefulness, and support for my projects: Dr. Andrew J. Taylor for guiding me through the entangled and intertwined chaos that is human physiology, Prof. Michael A. Cohen for patiently working me through the mathematics, analyses, and, at times, computer code, Drs. Anatoly Gorchetchnikov and Florian Raudies for brainstorming sessions and general advice.

My experience in Boston University's CNS and Math Departments enriched me both intellectually and personally primarily because of the outstanding people who work and study there, and who became friends. Thank you, Mark Kramer, Kyle Lepage, Eric Schwartz, Mukund Balasubramanian, Can Tan, Carol Jefferson, Cynthia Bradford, and many others for making it, well, fun! Without the help of Prof. Daniel Bullock and the financial support from CNS, this thesis would not have happened.

I offer my infinite gratitude to my family and friends who stood by me, sharing my ups and downs. My mom has been a source of inspiration, love, and perspective throughout my experience. My wife Nika patiently gave me space to write, stimulated me intellectually, and made the everyday enjoyable. Anya and Dima Granovski listened, listened, and listened again. Thank you!

**ADVANCES IN POINT PROCESS FILTERS
AND THEIR APPLICATION TO SYMPATHETIC NEURAL ACTIVITY**

(Order No.)

YEVGENIY ZAYDENS

Boston University Graduate School of Arts and Sciences, 2014

Major Professor: Uri T. Eden, Associate Professor of Mathematics and Statistics

ABSTRACT

This thesis is concerned with the development of techniques for analyzing the sequences of stereotypical electrical impulses within neurons known as spikes. Sequences of spikes, also called spike trains, transmit neural information; decoding them often provides details about the physiological processes generating the neural activity. Here, the statistical theory of event arrivals, called point processes, is applied to human muscle sympathetic spike trains, a peripheral nerve signal responsible for cardiovascular regulation. A novel technique that uses observed spike trains to dynamically derive information about the physiological processes generating them is also introduced.

Despite the emerging usage of individual spikes in the analysis of human muscle sympathetic nerve activity, the majority of studies in this field remain focused on bursts of activity at or below cardiac rhythm frequencies. Point process theory applied to multi-neuron spike trains captured both fast and slow spiking rhythms. First, analysis of high-frequency spiking patterns within cardiac cycles was performed and, surprisingly, revealed fibers with no cardiac rhythmicity. Modeling spikes as a function of average firing rates showed that individual nerves contribute substantially to the differences in the

sympathetic stressor response across experimental conditions. Subsequent investigation of low-frequency spiking identified two physiologically relevant frequency bands, and modeling spike trains as a function of hemodynamic variables uncovered complex associations between spiking activity and biophysical covariates at these two frequencies. For example, exercise-induced neural activation enhances the relationship of spikes to respiration but does not affect the extremely precise alignment of spikes to diastolic blood pressure.

Additionally, a novel method of utilizing point process observations to estimate an internal state process with partially linear dynamics was introduced. Separation of the linear components of the process model and reduction of the sampled space dimensionality improved the computational efficiency of the estimator. The method was tested on an established biophysical model by concurrently computing the dynamic electrical currents of a simulated neuron and estimating its conductance properties. Computational load reduction, improved accuracy, and applicability outside neuroscience establish the new technique as a valuable tool for decoding large dynamical systems with linear substructure and point process observations.

TABLE OF CONTENTS

ACKNOWLEDGMENTS	iv
ABSTRACT	v
TABLE OF CONTENTS	vii
LIST OF TABLES	x
LIST OF FIGURES	xi
LIST OF ABBREVIATIONS	xiii
1. INTRODUCTION	1
1.1 ORIGINAL CONTRIBUTION	3
1.2 THESIS OUTLINE	4
1.3 LITERATURE REVIEW	5
1.3.1 Muscle Sympathetic Nerve Activity	6
1.3.2 State-space dynamic estimation	10
2. MUSCLE SYMPATHETIC NERVE ACTIVITY CHARACTERIZATION	13
2.1 INTRODUCTION	13
2.2 METHODS	17
2.2.1 Spike Extraction	18
2.2.2 Subjects and Protocols	19
2.2.3 Visualization	21

2.2.3 Modeling Techniques.....	22
2.2.4 Statistics	24
2.3 RESULTS	24
2.3.1 Visualizations.....	25
2.3.2 Modeling.....	32
2.4 DISCUSSION	36
2.4.1 Why Spikes?	37
2.4.2 Mechanisms and Significance.....	40
2.4.3 Conclusions.....	44
3. RELATING MSNA TO HEMODYNAMIC VARIABLES	46
3.1 METHODS	50
3.1.1 Signal Processing.....	50
3.1.2 Visualizations.....	54
3.1.3 Spike Extraction.....	52
3.1.4 Subjects, Measurements, and Protocols.....	50
3.1.5 Modeling Techniques.....	54
3.1.6 Statistics	58
3.2 RESULTS	60
3.2.1 Visualizations.....	60
3.2.2 Modeling Results	66
3.3 DISCUSSION	73
4 MARGINALIZED PARTICLE FILTERS	78

4.1 METHODS	84
4.1.1 Regular Particle Filter	84
4.1.2 Marginalized Particle Filter	87
4.1.3 Marginalized Particle Filter with Point Process Observations	89
4.1.4 Algorithm	94
4.2 RESULTS	95
4.2.1 Model 1	96
4.2.1 Model 2	104
4.3 DISCUSSION	119
4.4 MPFPP EQUATIONS CHEATSHEET	122
5 CONCLUSIONS.....	125
5.1 Review of Thesis Results.....	126
5.2 Future Research Directions.....	128
BIBLIOGRAPHY.....	132
CURRICULUM VITAE.....	145

LIST OF TABLES

Table 1. Subject characteristics.....	19
Table 2. Subject summary statistics.....	20
Table 3. Kolmogorov-Smirnov statistics.....	33
Table 4. Magnitude model coefficients.....	69
Table 5. MPFPP Model 1 goodness-of-fit.....	100

LIST OF FIGURES

Figure 1. Muscle sympathetic nerve activity raster plot	26
Figure 2. Average spiking rates	27
Figure 3. Histogram of spike latencies	30
Figure 4. Spike latency densities	31
Figure 5. Density features as a function of rate	34
Figure 6. Power spectral density of spikes.....	61
Figure 7. Occupancy-normalized histogram.....	63
Figure 8. Coherence of spikes and hemodynamic variables.....	65
Figure 9. Preferred phase model parameters.....	66
Figure 10. Modulation magnitude model parameter.....	68
Figure 11. Nonparametric spiking rates.....	70
Figure 12. Proportion of spike train variance explained.....	72
Figure 13. Actual and estimated state variables.....	98
Figure 14. State estimation scatter plots	100
Figure 15. Joint linear/nonlinear posterior density	102
Figure 16. Root mean square error measures for Model 1	103
Figure 17. Latent variable estimates	111
Figure 18. Latent variable estimates (detailed).....	113
Figure 19. Maximum conductance estimates.....	114
Figure 20. Root mean square error for Model 2	115

Figure 21. Voltage and maximal conductance box plots..... 117

LIST OF ABBREVIATIONS

BP	Blood pressure
DAP	Diastolic arterial pressure
HF	Respiratory frequency (approx. 0.25 Hz)
IHE	Isometric handgrip exercise
LF	Mayer wave frequency (approx. 0.1 Hz)
MPF	Marginalized particle filter
MPFPP	Marginalized particle filter with point process observations
MSNA	Muscle sympathetic nerve activity
MSE	Mean square error
PFPP	Particle filter with point process observations
PSD	Power spectral density
R	Respiratory signal
RMSE	Root mean square error
SMC	Sequential Monte Carlo

1. INTRODUCTION

The analysis of the data recorded from nerve cells has been paramount in moving neuroscience forward. There are several methods to directly record electrical activity associated with neural information; one of the most prominent involves detection of cells' cross-membrane potential. We focus here on the problem of analyzing such recordings, specifically on interpreting sequences of stereotypical changes in cells' electrical activity; they are known as action potentials or spikes. Spike trains are usually obtained by the insertion or implantation into the nerve cell tissue of tiny electrodes that collect information about localized electrical activity. The recordings can either directly capture the voltage changes inside the cell (intracellular) or infer them from electric fields outside the cells (extracellular); this opens up access to the specific behavior of individual neurons as well as to the joint activity of neuronal ensembles. Analyzing these spike trains has helped study both micro- (e.g. intracellular signal propagation) and macro-level processes (e.g. rhythmic activities of clusters of neurons).

Analysis of the peripheral nerve recording in humans has been a challenge. First, the difficulty of accessing the peripheral nerves in vivo has limited the study to mostly cardiovascular and somatosensory neurons of the sympathetic nervous system. Getting to the nerves requires puncturing the skin with a relatively thick electrode, introducing the second problem: high levels of background noise. Analysis of these extracellular recordings, therefore, relied, and largely still does, on high degrees of averaging. This eliminated noise, but also masked vast amounts of potentially important details that are available in raw data. Only recently have there been introduced techniques that extracted

the individual cells' action potentials directly from the raw voltage trace and used them to study important features of the multifiber data (e.g. individual neuron contribution to the aggregate output of the bundle).

Both intra- and extracellular recordings shed light on how individual cells function. Because neurons process information with electric currents, modeling them is one of the best ways to describe the cells' functionality. Biophysical models, if constructed correctly, can predict with high degree of accuracy the cell's response to a variety of stimuli, from external (e.g. retinal cells' light response) to purely internal (e.g. decision-making, reflexive responses). The dynamics of cellular currents are complicated, and models describing them are complex as well. They often involve a significant number of both variables and parameters, and fitting these large-dimensional models to data becomes an extremely complex and computationally expensive task. In addition, finding suitable parameter values often is a tedious manual process. These complications make dimensionality reduction techniques extremely useful in computational neuroscience.

In this thesis we will utilize the spike trains obtained from the sympathetic nervous system in humans and study their properties without identifying individual neurons. We will use the spike train analysis framework to show that the changes in the properties of the multifiber spike trains are related to the average intensity of the nerve outflow. We will relate these changes to the concurrently collected physiological variables in an attempt to ground them in broader physiology. In addition, we will expand the spike train analysis framework by introducing a novel method of fitting

dynamic biophysical models to spike train data. Specifically, we will utilize the model's existing substructure to process its linear and nonlinear components separately; this will lead to increases in computational efficiency for large-dimensional models. In addition, we will show that this method can be used to estimate both the cell's electrical activity and the model parameters from only a set of spikes produced by the cell.

1.1 ORIGINAL CONTRIBUTION

The original contributions of this thesis are:

- a) An application of the spike train statistical analysis paradigm to the human sympathetic nerve activity.
- b) Identification of the structural changes in the multifiber sympathetic spiking as a function of the average nerve outflow and a rigorous statistical description of these changes.
- c) An assessment of the relationship between the multifiber sympathetic spiking patterns and those of the concurrently collected hemodynamic variables.
- d) Adaptation of the Rao-Blackwellization technique as described in Chapter 4 to the point process observations and the application of the resultant algorithm to large-scale biophysical models.

Materials related to the four articles above have been previously published in the following article

E. Zaydens, J. A. Taylor, M. A. Cohen, U. T. Eden, 2013, "Characterization and Modeling of Muscle Sympathetic Nerve Spiking," *Biomedical Engineering, IEEE Transactions on*, Vol. 60, No. 10, pp. 2914-2924

and a conference submission

E. Zaydens, C. O. Tan, J. A. Taylor, M. A. Cohen, U. T. Eden, 2011, "A novel approach to analyzing muscle sympathetic nerve activity in humans." *Cosyne Abstracts 2011*, Salt Lake City USA.

Chapters 3 and 4 are being currently prepared for publication as two separate journal articles.

1.2 THESIS OUTLINE

Current thesis is divided into five total chapters. The current chapter serves as a brief introduction to the overall topic of the work. It will also review the current literature on point process modeling in neurosciences and on the analysis of human muscle sympathetic nerve activity, focusing on modeling of its high- and low-frequency oscillations. Chapter 2 will describe our work on characterizing the high-frequency patterns of multifiber sympathetic nerve activity and relating these patterns to spike trains' low-frequency oscillations. Chapter 3 will expand this analysis to include hemodynamic variables. Specifically, we will visualize and model the relationships between hemodynamic covariates and the sympathetic spiking in two low-frequency bands that were previously determined to have physiological relevance. We will compare how well different covariates predict spiking within these bands and draw physiological conclusions. Chapter 4 will expand the spike train analysis theory to include marginalization into the state-space estimation paradigm that uses point process observations. It will outline the "curse of dimensionality" problem and will exploit a model's linear and Gaussian substructure to solve it. The proposed technique will be used to estimate parameters of the well-known model of nerve cell internal currents. Chapter 5 will summarize the work and draw some inferences. It will also propose

avenues for future work, including expanding the range of applicable physiological problems as well as refinement and customization of the proposed methods.

1.3 LITERATURE REVIEW

Among techniques to collect peripheral nervous system data in humans, microneurography provides the most detailed account of the neuronal activity. Microneurography can be used to collect several different types of data, including afferent sensory activity (e.g. muscle spindle impulses, discharges from the tendon organs, muscle and skin nociceptors as well as skin mechanoreceptors) and efferent control activity (e.g. skin hair motor and renal activation). Because we are interested in systemic peripheral effects, this review will focus on the efferent muscle sympathetic nerve activity (MSNA), which regulates peripheral blood flow by managing contractile forces of the smooth muscle comprising the resistance vessel walls. Specifically, we will review the history of MSNA collection, the analysis tools used to process the collected data, and conclusions thereby derived about the physiology of high- and low-frequency content of the nerve outflow. We will focus predominantly on the analysis of human MSNA.

In this section we will also cover literature describing state space estimation methods that track the temporal trajectory of a process by observing only its output. Different state evolution and output production models can be constructed to accommodate various state trajectories and observation types. We will start by covering the simplest designs that place the most restrictive assumptions on the process and observation models. We will then move on to techniques that relax these assumptions and

therefore can tackle substantially more difficult estimation problems. Specifically, we will first focus on the famous Kalman filter and its extensions, then cover the Sequential Monte Carlo (SMC) technique, and end by describing the SMC improvements that take advantage of the models' substructure to improve efficiency and accuracy of the estimation.

1.3.1 Muscle Sympathetic Nerve Activity

The first reported microneurographic recording of muscle sympathetic nerve activity is due by Hagbarth and Vallbo in 1968. Without using anesthesia they inserted tungsten microelectrodes through skin directly into the nerve bundle containing sympathetic axons. They found that the combined electrical sympathetic activity was synchronized to the cardiac cycle of the subject and also affected by subject's respiration. In 1971 Delius, Hagbarth, Hongell, and Wallin [26] used a sympathetic ganglion blocking agent to abolish such voltage traces, confirming that these are indeed postganglionic muscle nerve signals. They also showed that MSNA activity was inversely correlated with the blood pressure signal, demonstrating the vasoconstricting effects of MSNA activation. In a subsequent study [25] they tracked how the MSNA changed during various conditions. These experiments showed that the baroreflex was the predominant mechanism regulating the sympathetic outflow, but that other reflexes were taking on partial control during some experimental maneuvers.

Grassi and Esler [47] describe ways of measuring the sympathetic activity in humans. Microneurographic measurements are generally expressed in number of bursts per minute and, correcting for different heart rates across subjects, number of bursts per

100 heart beats. Burst size information (burst height and total area) is also collected, as is burst timing. Fagius, Sundlof, and Wallin [40] have measured the latency of the individual bursts during different conditions, and found it to correspond to the conduction delay they measured directly [41]. Bursts latency continues to be the most detailed temporal information that is possible to extract from the burst-based measurements.

Several groups increased the impedance of the microelectrode and recorded individual action potentials directly. These spikes were then used to study the firing properties of the individual fibers. Macefield and colleagues found that individual neurons mostly fire once per cardiac cycle [82] and that their firing probability plays a role in the acute growth of sympathetic outflow in response to various stressors [83]. They also found positive correlation between the latency of individual firings and the diastolic pressure of the cardiac cycle that generated the spike, suggesting that fibers with higher conductance velocities are turned on by dips in blood pressure [81]. Tsukahara and Mano confirmed the pressure-latency link by computing fibers' firing probability: fibers with both higher threshold blood pressure and longer spike latency were found to have higher firing probability [127]. This corroborated the findings in [81] that faster-conducting fibers were engaged by dips in blood pressure, but that this seldom happened. Murai et al. expanded the single-fiber work by detecting differences in how individual fibers respond to different stressors. Specifically, they found that multiple firings of individual fibers play a larger role in the growth of MSNA outflow caused by Valsalva maneuver than by IHE [92].

Microneurography with increased-impedance electrodes is a complex procedure that can only detect small number of individual fibers. To circumventing this limitation, Salmanpour and colleagues used wavelets to extract individual action potentials from the multifiber signal obtained using standard microneurography technique [109]. They confirmed earlier findings that different type of sympathoexcitatory stressors elicit different fiber recruitment patterns. By classifying the amplitudes of action potentials, they showed that no new amplitude classes activated in response to lower-body negative pressure [110], but that large action potentials started appearing during sympathoexcitation induced by end-respiratory breath hold [109]. They also found that different action potential sizes seem to be controlled by different physiological mechanisms: small- and medium-size action potentials activated readily in response to baroreflex, but large action potentials exhibited lack of response [111].

Human MSNA exhibits oscillations at two low-frequency bands: Mayer wave frequency of about 0.1 Hz and respiratory frequency that usually falls in the [0.2, 0.3] Hz range. Two theories of Mayer wave origination have been put forth. The first says that they are produced by a pacemaker residing in the central nervous system. The evidence for this theory comes from animal studies. Kaminski et al. found that slow waves in lumbar SNA in vagotomized dog can be induced by raising the subarachnoid pressure [71], suggesting that sympathetic oscillator can be located in the spinal cord. Preiss and Polosa detected slow fluctuations in preganglionic sympathetic activity, placing any possible oscillator into a central nervous system [99]. Using spectral analysis, Montano

and colleagues found a low-frequency oscillations in some medullary sympathetic-related neurons in cats [91], corroborating a central nervous system oscillator theory.

The second theory states that Mayer waves constitute feedback resonances of the baroreflex loop. In particular, the filtering function of the loop is thought to possess a phase shift that introduces an in-phase input-output relationship, producing resonant oscillations [67]. The evidence for this theory also comes mainly from animal studies. Baroreflex loop was opened and its transfer function analyzed in dogs [114, 78], rabbits [61], and rats [112]. In all these studies, the resonant frequency was found to be close to the Mayer wave frequencies of the species.

Several mechanisms have been proposed to explain MSNA variability at the respiratory frequency band. Mano et al. [85] identified four potential ones: central oscillatory theory, whereby a single central oscillator is responsible for initiation of both respiratory muscle contractions and MSNA efferent outflow. Second is alveolar inflation reflex theory, which states that alveoli expansion causes changes in MSNA. Third, there is a baro- and cardiopulmonary reflex theory that links respiratory MSNA fluctuations to those of arterial pressure. And finally, there is a chemoreflex theory, whereby respiratory changes in blood gas levels affect MSNA. Evidence for some of these is available from human subjects. St. Croix et al. confirmed that alveolar expansion affects MSNA by capturing an increase in nervous outflow during passive positive pressure ventilation [121]. Ikeda et al. 1997 confirmed these findings by showing that MSNA is increased independent of the arterial pressure during spontaneous positive-pressure breathing [60]. Somers et al. provided evidence for the chemoreceptive reflex theory by showing that

MSNA surged with no change in arterial pressure during last 5 seconds of Mueller maneuver [120]. St. Croix and colleagues also showed that baroreflex remains active during all experimental maneuvers [121], suggesting that noninvasive experiments are likely to fail in completely separating reflexes affecting MSNA at respiratory frequencies.

1.3.2 State-space dynamic estimation

Rudolf E. Kalman along with Richard S. Bucy have developed what is now known as Kalman-Bucy (or simply Kalman) filter in the beginning of 1960s [69, 70]. Kalman filter provides an analytic solution to a problem of dynamic estimation of latent processes from the observations they produce. Kalman filter assumes that a) both the process and its observable output are Gaussian random variables, b) state evolution model is linear, and c) observations linearly depend on the state variable. Because engineering problems are seldom linear and are often non-Gaussian, an extension was immediately developed that linearized a nonlinear state transition and observation functions about the current state estimate [44]. This Extended Kalman filter allowed working with nonlinear models, but diverged if the model was misspecified or if the initial estimate was wrong.

To address the inability on the Kalman-type filters to deal with highly nonlinear data, a methods that put no *a priori* constraints on the model was developed by J. E. Handschin and colleagues in 1969 [55, 54]. They approximated the distribution of a state process given the observed data (posterior probability) with a set of random weighted samples. This method, known as Sequential Monte Carlo or particle filter, was computationally intensive and suffered from a sample degeneracy problem, where

weights of all but very few particles became negligible after several estimation steps. The sample degeneracy was addressed by Gordon et al. in 1993, who devised a way to rejuvenate the samples by making copies of the high-weight samples and discarding the low-weight samples called resampling [46]. Resampling improved the algorithm's stability, and, along with the increase in available computational power, led to its wide acceptance.

Sequential Monte Carlo estimator was improved by several groups. Pitt and Shephard proposed favoring the particles that were likely to survive at the next time step as opposed to resampling based strictly on the particle weights [97]. Doucet et al. unified several previously proposed resampling methods into a probabilistic framework simplifying theoretical analysis of the algorithms [32]. Godsill and colleagues introduced particle smoothing, whereby the estimate at a specific time is improved by using future "look-ahead" observations [45]. Schön et al. improved the accuracy and the efficiency of particle filtering by separating the linear substructure of the state transition model and estimating this conditional linear part with a random mixture of Gaussians [116].

Because of its lack of intrinsic *a priori* modeling assumptions particle filters were embraced by the neuroscience community, that often faces non-linear and non-Gaussian neural data. For instance, Brockwell et al. used a particle filter to recover hand's position, velocity, and acceleration traces from spike trains recorded in a premotor brain area of a rhesus monkey [13]. Wang et al. improved on this procedure by introducing additional observation channels, such as somatosensory cortical spike trains [134]. Ergun and colleagues employed a particle filter to decode the temporal changes in the receptive field

of rat's hippocampal neurons as the animal foraged in an open environment [37]. By incorporating model parameters into the state vector of the particle filter, Meng et al. concurrently estimated both the voltage dynamics and the biophysical variables (e.g. maximal conductance) of a neuron using only the times when the neuron spiked [89].

2. MUSCLE SYMPATHETIC NERVE ACTIVITY CHARACTERIZATION

Muscle sympathetic nerve activity (MSNA) is a primary source of cardiovascular control in humans. Traditional analyses smooth away the fine temporal structure of the sympathetic recordings, limiting our understanding of sympathetic activation mechanisms. We use multifiber spike trains extracted from standard microneurography voltage trace to characterize the sympathetic spiking at rest and during sympathoexcitation. Our analysis corroborates known features of sympathetic activity, such as bursting behavior, cardiac rhythmicity, and long conduction delays. It also elucidates new features such as large heartbeat-to-heartbeat variability of firing rates and precise pattern of spiking within cardiac cycles. We find that at low firing rates spikes occur uniformly throughout the cardiac cycle, but at higher rates they tend to cluster in bursts around a particular latency. This latency shortens and the clusters tighten as the firing rates grow. Sympathoexcitation increases firing rates and shifts the burst latency later. Negative rate/latency correlation and the sympathoexcitatory shift suggest that spike production of the individual fibers contributes significantly to the control of the sympathetic bursts strength. Access to fine scale temporal information, more physiologically accurate description of nerve activity, and new hypotheses about the nervous outflow control establishes sympathetic spiking as a valuable tool for the cardiovascular research.

2.1 INTRODUCTION

Because of its importance in maintaining cardiovascular homeostasis, efferent muscle sympathetic nerve activity in humans has been an active area of research for

several decades [8]. It is traditionally assessed by collecting raw voltage traces with an electrode subcutaneously inserted into a peroneal nerve, and converting these traces into temporally smoothed signals with clearly identifiable bursts of sympathetic activity [53]. The bursts are quantified using the following measures: strength (e.g. burst amplitude or burst area), frequency (e.g. number of bursts per minute or per 100 heartbeats), and crude timing information (e.g. burst latencies). These measures have been successful in identifying human sympathetic activity patterns across reflex activations such as baroreflex and metaboreflex [128, 131, 39] as well as examining characteristic changes in outflow with pathological conditions such as hypertension, congestive heart failure, and obstructive sleep apnea [19, 38, 42].

Animal models reveal that high frequency content of the sympathetic activity can substantially influence vascular responses. For instance, the constriction of resistance vessels is augmented by the intra-burst pulse frequencies during stimulation with fixed inter-burst frequencies; the maximal response is attained at 10 Hz in rabbits [4] and at 40 Hz in cats [3]. Also, the perfusion pressure was strongly correlated with the carrier wave frequency of up to 20 Hz during a frequency modulated stimulation of the sympathetic nerve trunk in the dog hind limb [100]. The smoothing inherent in the traditional MSNA analysis technique leads to a substantial loss of temporal resolution in the recorded sympathetic activity, leaving it unable to account for such high-frequency sympathetic activity content.

Augmenting the microneurography setup by increasing the impedance of the electrodes allowed researchers to record some of the action potentials emitted by the

individual sympathetic fibers, thus capturing a portion of high-frequency structure present in the sympathetic output. More importantly, the analysis of single-fiber firings identified the causes of the inter-individual and cross-conditional burst size variability for a large number of conditions [81, 92, 127, 36, 83]. Increasing the electrode impedance to meticulously track the firings of single fibers considerably limits the number of axons that are sensed. For instance, only 14 total units from 8 subjects were recorded by Macefield et al.[82], 34 total units from 16 subjects were recorded by Elam et al.[36], and 39 total units from 5 subjects were recorded by Tsukahara and Mano [127]. This is compared with 1100 sympathetic axons found in the rat peroneal nerve[115]. The number in humans is likely close: the lateral fascicles of deep peroneal nerve innervating muscle and not skin were found to contain a minimum of 195 (S.E.=18) myelinated fibers [123], and mammalian peroneal nerve is known to have unmyelinated C to myelinated A fiber ratio >1 [66, 98]. Accessing such a small subset of available axons makes it hard to assess in humans the complex patterns of sympathetic nerve firing exhibited by animal models (e.g. only 57% of rat muscle sympathetic neurons responded to nociception stimulus [51], 8% of such neurons had no respiratory modulations[52]; a detection in cat of muscle sympathetic efferents with low cardiac rhythmicity [9]).

A new technique was designed that addressed the limitation on the number of recorded fibers: the extraction of individual spike information directly from the raw multifiber voltage trace obtained using standard-impedance microneurography [29, 17, 124]. A large number of such studies identified the action potential composition of the multifiber sympathetic bursts obtained with the traditional MSNA analysis [122, 12, 110,

86]. To do so, individual action potentials were extracted and classified based on their shape, obtaining up to 24 morphology types per subject [122]. Focusing on the burst structure left parts of the spike train unutilized, however. For instance, Steinback et al. [122] and Salmanpour & Shoemaker [111] extracted spikes from the whole recording, but only used the ones that coincided with the bursts. Maslov et al. [86] only extracted spikes from parts of the recording that formed the multifiber bursts. In addition, a full characterization of the statistical properties of the MSNA spiking patterns is, to our knowledge, currently missing. For example, the physiologically important high-frequency multifiber firing within the MSNA bursts (e.g. spike placement variability) remains unexplored, as does its relationship to the low-frequency spiking structure (e.g. variability in burst sizes). In fact, most multifiber studies do not use the spikes themselves. For instance, Brychta et al. [16, 17] smoothed the spike train to analyze MSNA spectral signature and to model low-frequency arterial pressure oscillations. Also, while inhomogeneous Poisson models have been successfully applied to the analysis of spiking structure in a number of neural systems, we are not aware of them being utilized to identify specific spike patterns in multifiber sympathetic efferent activity. Moreover, statistical tools ignoring the precise spike timing (a critical information carrier for many neural systems [118, 20, 7, 48]) are often favored. For example, Greenwood et al. [48], Murai et al. [92], and Brychta et al. [17] first inferred firing rates and then applied standard statistical methodology (e.g. least squares regression) to test hypotheses and fit models. Therefore, the information currently present in the multifiber spike trains is not fully exploited. In particular, we do not know how the placement of spikes changes as a

function of time from the associated R-peak; we do not know if such placement is different for large vs. small bursts; we do not know if spikes occur between bursts.

We address some of these questions in this paper, which has two main goals. Methodologically, we aim to combine in a novel way the existing standard-impedance spike extraction algorithm with a successful modeling framework in order to extract and utilize more information from the raw voltage trace than was previously possible. Scientifically, we aim to describe the aggregate firing patterns of the multifiber sympathetic spike trains (including firing within regions traditionally considered silent) and examine in detail the interactions between their high and low frequency components.

Our analysis is structured as follows: first, the raw voltage trace is denoised and sequences of action potentials are extracted. The temporal structure of the resultant multifiber spike trains is visualized using rasters plots. A point process model is applied to assess the distribution of high frequency spikes as a function of the average spike rate. The model is fit to the spike trains observed at rest and during a sympathoexcitatory maneuver. Maximum likelihood estimates of the model parameters are compared across subjects. Finally, the similarities and differences in spiking are examined in the context of known physiology.

2.2 METHODS

This section covers the methods used in this study. We first detail the data collection and spike extraction process as well as subject selection criteria. We then describe the visualizations we used to help us arrive at the model we eventually utilized, which is described last.

2.2.1 Spike Extraction

A sympathetic activity assessment technique was devised [124] that discriminates neural firing from physiologically-irrelevant information, and which accurately identifies multifiber nerve activity independent of any *a priori* criteria. This method is based on the statistical characteristics of the different artifacts, noise, and multifiber nerve activity that compose the unfiltered recording. A variant of the Expectation-Maximization algorithm was used at successive stages of processing to separate these components to extract the multifiber spike train. The details of the algorithm are described in [124], and, briefly, consist of the following steps: line noise and muscle twitches are removed from raw voltage recordings, producing a set of action potentials. These are then treated as a convolution of a set of pulses (Kronecker delta functions) with a prototypical spike template, which is extracted from the data. The posterior probability of a spike given the data is computed at each sample by deconvolving the spike template from the signal and removing the background noise using a mixture of distributions model. The posterior probabilities are then thresholded and binarized, producing a spike train at 40 kHz. Note that by using the action potential template we assumed that the shape of a subset of action potentials in the signal is roughly similar and that some fraction of the action potentials is considerably larger in peak amplitude than the noise source. Beyond that, both spike shape and amplitude are explicitly ignored because they are heavily influenced by the distance from the axon to the electrode and are highly variable across recording sessions. We did not sort the spikes because we were aiming to characterize how the spike patterns produced by *all* active fibers change as a function of outflow intensity and were not

looking to study how different fiber types influenced the pathway activation responses. The output of the spike extraction algorithm, therefore, is a binary multifiber spike train where neuronal firings by different axons cannot be differentiated. A refractory period was enforced by removing any of the spikes within 0.5 msec of each other. Average morphology templates were then collected and compared to noise samples and prototypical action potentials to ensure that spike locations in the binary train mark voltage deviations that are statistically significant and resemble actual spikes. We used this technique to extract multifiber spike trains from eight subjects under baseline and sympathoexcitatory conditions. Baseline activity was collected during supine rest, and an isometric handgrip exercise (IHE) served as a sympathoexcitatory stimulus.

Age (yr)	Height (m)	Weight (kg)	BMI (kg/m²)
25.9, [21, 30]	1.71, [1.52, 1.88]	62.9, [47.6, 75.1]	21.7, [16.4, 32.5]

Table 1. Subject characteristics: cross-subject means and ranges for age, height, weight, and body-mass index (BMI).

2.2.2 Subjects and Protocols

All volunteers gave their written informed consent to participate. All protocols were approved by either the ethical committee at the Grenoble University Hospital Center or institutional review board at Spaulding Rehabilitation Hospital, and conformed to the Declaration of Helsinki.

We studied healthy men (n=4) and women (n=4) with no signs or symptoms of hypertension, diabetes, cardiovascular or neurological disease or cancer, with normal resting ECG, and no regular use of tobacco. The population statistics are available in Table 1.

	Baseline		IHE	
	Mean	Range	Mean	Range
Duration (sec)	414.5	[300.6, 609.8]	264.8	[186.6, 387.7]
RR lengths (sec)	1.04	[0.70, 1.58]	0.79*	[0.52, 1.49]
Mean Rate (Hz)	7.30	[3.51, 12.54]	14.89*	[8.74, 26.96]
Mean Interval Latency (sec)	1.27	[1.07, 1.46]	1.29	[1.11, 1.45]
Subject Number	Average firing rate (Hz)	Fano Factor	Average firing rate (Hz)	Fano Factor
1	12.30, [0, 65.09]	7.55	10.81, [0, 66.57]	4.71
2	11.16, [0, 101.58]	10.10	22.71, [0, 98.07]	12.48
3	15.36, [0, 55.78]	6.40	18.87, [0, 68.36]	7.79
4	7.23, [0, 52.33]	5.30	17.02, [0, 273.26]	31.51
5	11.51, [0, 77.61]	5.70	11.65, [0, 57.65]	5.05
6	5.59, [0, 36.1]	4.08	11.02, [0, 40.107]	5.65
7	9.86, [0, 162.99]	36.27	32.46, [0, 244.48]	32.87
8	4.92, [0, 23.95]	1.99	23.79, [0, 267.46]	47.95

Table 2. Top panel: summary statistics for the eight subjects in the study. Includes cross-subject means and ranges of the following: recording durations, mean duration of RR intervals within a recording, mean spiking rate within a recording, and mean latency of spikes within the second period of elevated firing. Bottom panel: heartbeat-to-heartbeat variability of cardiac cycle firing for individual subjects expressed as the Fano factor, a ratio of variance of cardiac cycle spike counts to their mean.

*indicates statistically different from baseline ($p < 0.05$).

Baseline activity data was derived from a five-minute period of supine rest. Sympathoexcitatory data was obtained by subjects performing isometric handgrip exercise at 35% maximal voluntary contraction (MVC) force sustained to fatigue. (The subject had previously performed three MVCs on a handgrip dynamometer to determine 35% MVC.) During the exercise, target force was displayed on a computer monitor. Subjects had continuous visual and auditory feedback from the investigators to ensure maintenance of target force until exhaustion, which was defined by the inability to maintain a target force for more than 3 sec despite continued effort and verbal encouragement [10].

2.2.3 Visualization

Extracting individual spikes allowed us to analyze the sympathetic activity using visualization techniques designed specifically for spiking data. We constructed raster plots, which display spikes as individual dots in time on the x-axis, with multiple trials or time intervals stacked on top of each other on the y-axis. In our case, the intervals are triggered by the ECG R-peaks, and the spiking activity for the following two-second period is displayed as a function of time elapsed from the R-peak (latency). Each two-second interval includes the spiking activity over a minimum of two cardiac cycles following each R-peak (possibly more for shorter RR intervals during IHE).

Based on these visualizations, and on the known conduction delays for sympathetic firing [40, 130], we focused on an interval of high-density spiking from 700-1400 ms after each R-peak. We computed the average latency within each interval as a sum of latencies of spikes within the interval divided by its number of spikes. We also computed an average firing rate for each interval as the number of spikes in an interval divided by its duration. We visualized the distribution of these average firing rate values over heartbeats using box plots, histograms, and by directly plotting the average rate time series. To assess the heartbeat-to-heartbeat variability of the sympathetic outflow, we computed the Fano factor, a ratio of the variance of cardiac cycle spike counts to their mean. The Fano factor has a theoretical value of one for Poisson processes.

In order to visualize the patterning of spiking activity as a function of the average firing rate, we modified the raster plots by sorting the intervals on the y-axis from those with the highest average rate to those with the lowest average rate. We also examined two

subsets of intervals: those with low average firing rates below 10 Hz, and those with high firing rates above 40 Hz, and computed for each a smoothed histogram of the spiking as a function of latency. This was accomplished by convolving the histogrammed spikes with a boxcar kernel with a width of 10 msec.

2.2.3 Modeling Techniques

We constructed a probability model to describe a pattern of spiking activity of the multifiber sympathetic nerve bundle as a function of the average firing rate over each interval of high density spiking. Specifically, we binned the spiking data (1 msec bin width) and modeled it as an inhomogeneous Poisson process conditioned on the average firing rate within each interval. At any given rate, the model describes the probability distribution of the latency relative to the associated R-peak, at which the spikes are likely to occur. Mathematically, let R be the average firing rate computed in a given interval of high-density spiking, let τ be the latency relative to the associated R-peak, and let p be the probability that a randomly chosen spike within the interval occurs at latency τ , given that the interval has an average firing rate of R . Then, p is related to the firing rate function $\lambda(\tau)$ of the spiking process according to the formula: $p(\tau) = \lambda(\tau) / \int_{\tau_start}^{\tau_end} \lambda(u) du$, where τ_start and τ_end are starting and ending latencies of the interval. We fit the following model of the spike latency probability p :

$$p(\tau) = C \cdot \exp \left\{ \begin{array}{l} \beta_1 \tau + \beta_2 \tau^2 + \\ + \beta_3 \tau R + \beta_4 \tau^2 R + \beta_5 \tau^3 R + \\ + \beta_6 \tau R^2 + \beta_7 \tau^2 R^2 + \beta_8 \tau^3 R^2 \end{array} \right\}. \quad (1)$$

Here C is a normalization constant that ensures that p integrates to 1, and the exponential

function guarantees that the probability density is positive. The model parameters $\beta=(\beta_1,\beta_2,\beta_3,\beta_4,\beta_5,\beta_6,\beta_7,\beta_8)$ are fit by maximum likelihood, i.e. they are selected to maximize the probability of observing the data.

For each interval of high density spiking in the recording, we plug its average firing rate R into the model to obtain a probability distribution function for the location τ of any arbitrary spike. We compute the mode of this probability distribution to find the expected latency of the spiking relative to the associated R-peak, and we compute the variance of this distribution to describe the concentration of spiking about the expected latency. We also compute the maximum instantaneous rate reached within an interval, which defines the peak latency and describes the maximum output of the spike train in spikes/sec.

Along with maximum likelihood estimators of the model parameters β , we compute their variance-covariance matrix. Using that information, we draw 100,000 samples from the asymptotic distribution of β . We use these drawn values to approximate the sampling distribution of the expected latency, the spike concentration, and the maximum instantaneous rate as well as to construct their 95% confidence intervals.

The latency range for each condition was computed by subtracting the expected latency for the burst with the highest rate from the expected latency of the burst with the rate of 10 Hz, which is the slowest spiking rate at which latency is reliably computed. Latency ranges were then averaged across subjects for baseline and IHE to obtain a cross-subject mean latency range for each condition.

2.2.4 Statistics

Simple linear regression was performed to compute linear changes in average spike rates over time and the significance of the estimated slope parameters. The significance of the parameter estimates in Eq. 1 was calculated based on the observed Fisher information computed from the data. Paired t test was used to compare RR lengths, mean rates, and mean interval latency across conditions. P-value <0.05 was considered significant. Kolmogorov-Smirnov tests that compared the modeled spike latency distribution p (Eq. 1) and the empirical one were used to assess the model goodness-of-fit for each subject/condition.

2.3 RESULTS

Table 2 displays the spiking statistics for all subjects in this study. The recording duration varied from subject to subject, was generally longer at rest, and reached a maximum of about 10 minutes. As expected, the length of an average RR interval decreased from 1.04 seconds during baseline to 0.79 seconds during IHE, reflecting increased heart rate across all subjects during isometric handgrip exercise. The spiking rates were different for each subject and varied about three-fold within each condition. During isometric exercise, the mean spiking rate averaged across all subjects was 14.8 Hz, about twice that at rest. Fano factors were highly variable across subjects, with the minimum value of 1.99 for all recordings in the study. All values were significantly larger than 1 (p-values $<10^{-5}$ for all subjects), showing higher heartbeat-to-heartbeat variability in the firing rate than would be observed for Poisson Processes. The cross-subject means of average spike latencies for both baseline and IHE were 1.27 and 1.29

seconds (see Table 2). The means were not statistically different. The cross-subject mean latency range obtained with the spiking model was 79 (S.E.=20) msec for baseline and 86 (S.E.=19) for handgrip.

2.3.1 Visualizations

Figure 1 depicts the raster plots for a representative subject under baseline (left, black) and IHE (right, gray) conditions and highlights the aggregate characteristics of the sympathetic spiking by marking the latencies of all spikes trailing each R-peak by less than 2 seconds. The most visually striking feature is the presence of the three regions of increased sympathetic firing in a two-second interval following every R-peak: one at a latency of about 300 msec, one at 1100 msec, and the third at about 2 seconds. The latency of the second high-density area (700-1400 msec) matches the average conduction delay described in the literature [131]. The reported progressive increases in the sympathetic outflow and the decreases in the RR intervals during IHE [72] do not affect the latency of the spikes in the second area of high activity. This is observed by comparing the values of the three high-density areas in the right plot of Fig. 1: the latency of the middle one stays constant at around 1100 msec for the duration of the recording, but the latencies of the first and the third areas have shifted at the conclusion of the IHE experiment. In addition, the shortening of the average RR intervals and the increase in the average spiking rates (see Table 2) contributed to the higher density of spikes during IHE; notice that the individual spikes are distinguishable in the beginning of the IHE recording in Fig. 1, but blur together towards its end. These features were observed in all subjects. The only cross-subject variation was the timing of the high-density spiking from

their corresponding R-peaks; the ranges of the mean spike latency for the second high-density region are found in Table 2.

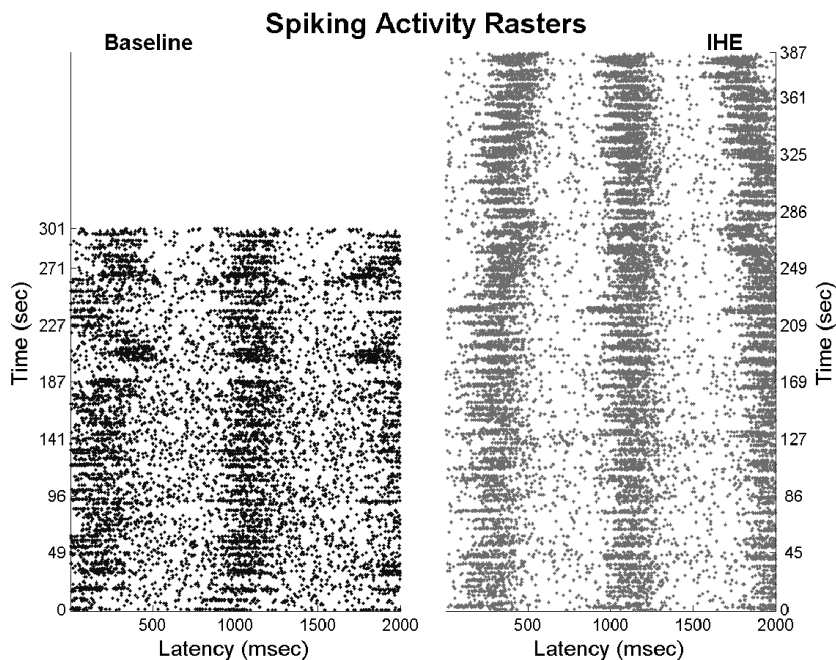


Figure 1. A raster plot of MSNA spiking activity triggered by ECG R-peaks for baseline (left, black) and IHE (right, gray) conditions. Each row illustrates 2 seconds of the spiking activity following an R-peak. The x-axis represents the time elapsed since the corresponding R-peak. The y-axis depicts the time of the R-peak corresponding to that row. Note that one particular spike might appear on multiple rows at various latencies relative to different R-peaks. For each heartbeat, the spiking activity increases 900-1300 msec after the R-peak, consistent with conduction delays reported in the literature. This latency is consistent across baseline and IHE conditions, despite the increasing heart rate during IHE. The intensity of spiking progressively increases during IHE. "Horizontal streaking" in the observed spiking activity is indicative of the heartbeats with reduced activity followed by heartbeats with elevated activity.

Another salient characteristic of the sympathetic spiking was the white spaces cutting through the columns of high activity (see Fig. 1). These horizontal streaks represent heartbeats with limited or no sympathetic firing interspersed with the cardiac cycles exhibiting high density of spikes. Such changes in the number of sympathetic spikes from one heartbeat to the next can be thought of as beat-to-beat variability of spiking, a quantity that is commonly measured by the Fano factor. Table 2 shows the

sample Fano factor values ranging from 1.99 to 47.95 across all subjects/conditions; they are considerably larger than the theoretical value of one that Fano factor attains for an inhomogeneous Poisson process. This large discrepancy indicates that heartbeat-to-heartbeat variability of the spike counts is significantly larger than what an inhomogeneous Poisson process would predict.

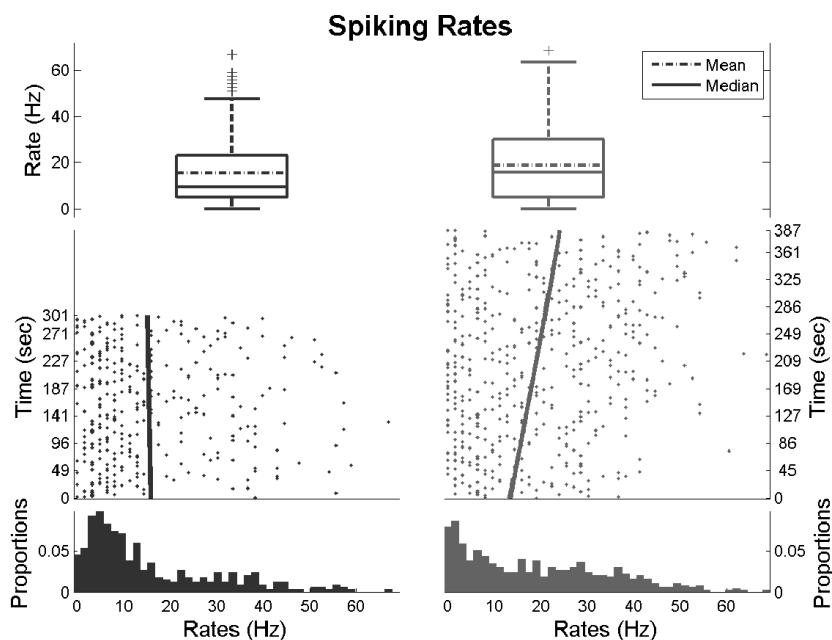


Figure 2. Average spiking rates in baseline (left, black) and IHE (right, gray) conditions. Top panel: box plots of spiking rates across heartbeats. Median and upper whisker of the firing rate are increased in IHE. Middle panel: scatter plot of firing rate on each heartbeat with linear regression fit. Firing rate is stable in baseline condition and increasing in IHE. Bottom panel: histogram of firing rates. Firing rates range from 0 to about 65 Hz in both conditions, but IHE has increased beat-to-beat variability relative to baseline; more heartbeats have virtually no spiking, and more heartbeats exhibit high rates of firing. Linear regression equation is $R=16.0183-0.0043t$ for baseline and $R=13.4179+0.0273t$ for IHE.

The intensity of the sympathetic activity for each recording was assessed by analyzing the average firing rates within intervals of elevated firing occurring at latencies of 700 to 1400 msec. The following were constructed to describe the rates: box plots,

time series diagrams, and histograms of the rate distribution. The top panel of Figure 2 displays box plots that describe the quartiles of the average rate distribution for a representative subject under baseline (left, black) and IHE (right, gray) conditions. Note that the top whisker in baseline and IHE conditions is large when compared to both the size of the box and of the lower whisker. This indicates that the range of the rates contained in the top quartile is much higher than that of the rates contained in the lower three quartiles. Also note that the maximum rate attained for both recordings is much higher than the upper boundary of the box, indicating that the rates in the top quartile of the distribution are appreciably higher than in the rest of the quartiles. In fact, the rate distributions for all subjects demonstrate similar patterns: the maximum rate for some recordings was 7-15 times higher than the 75th percentile and 9-16 times higher than the mean. Note, however, that while there was a clear increase in sympathetic firing in IHE when compared to baseline, the maximum firing rates were not considerably different between baseline and IHE for 5 out of 8 subjects (see Table 2). For instance, while subject 4 had IHE maximum rate higher than baseline, subject 5 had this relationship reversed.

The diagrams in the middle panel of Figure 2 revealed the time-dependent features of the spiking. First, average rates are increasing as a function of time during handgrip exercise (for a subject in Fig. 2 the IHE slope is 0.0273 Hz/sec, significantly larger than zero at $p=4.14 \cdot 10^{-6}$). This upward trend was found in all subjects; the average slope was 0.0609 (S.E.=0.0257) Hz/sec, significantly larger than zero ($p=0.0248$). Second, there is a large variability in the rates of the neighboring regions. For instance,

the highest rate in Figure 2 baseline is 66.9 Hz; the region exhibiting this rate is surrounded by regions with significantly lower rates of 12.7 and 3.2 Hz. Third, the amount of high- relative to low-rate regions stays constant for baseline but slightly increases for IHE (note a large quantity of the low-rate regions still present at the end of IHE, where the expected rate is highest). In addition, sympathoexcitation increases the rates within the high-frequency bursts of sympathetic spiking. For example, the mean rate for the regions with rates over 10 Hz occurring in the last 30 seconds of IHE was 39.64 (S.E.=5.96) Hz, which was significantly larger than the mean rate for similarly high-rate regions in the first 30 seconds of IHE (25.94, S.E.=1.51 Hz, $p=0.0257$) and in the last 30 seconds of baseline (22.75, S.E.=2.24 Hz, $p=0.0095$).

In order to directly compare the outflow strength at rest with that during peak sympathoexcitatory response, we compared mean rates in the first and last 30 seconds of IHE. A large increase was found in all subjects, the cross-subject means are 8.60 (S.E.=0.96) Hz in the first 30 seconds of IHE to 21.94 (S.E.=4.85) Hz in the last. To ensure that the difference is not an artifact of the time window placement within a recording, we compared the mean rates in the last 30 seconds of baseline to the last 30 seconds of IHE; an increase was found in all subjects with means of 6.83 (S.E.=0.74) Hz vs. 21.94 (S.E.=4.85) Hz (statistically significant).

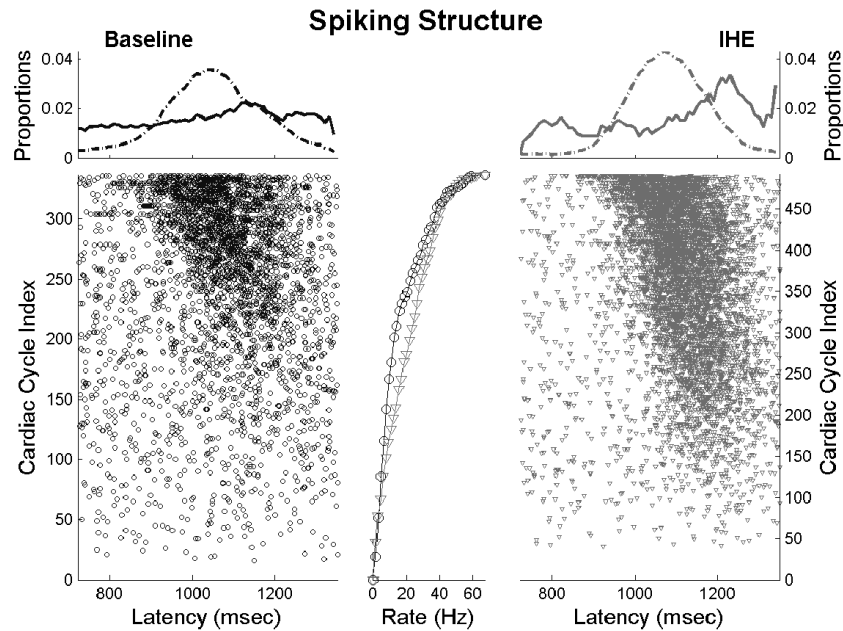


Figure 3. Top panel: histogram of spike latencies for low spiking (<10 Hz, solid) and high spiking (>40 Hz, dot-dashed) regions for baseline (left, black) and IHE (right, gray) conditions. Low-rate spiking occurs uniformly across all latencies shown, while high-rate spiking tends to concentrate about a preferred latency with a bell-shaped distribution for both conditions. Bottom panel: raster plots of the MSNA spiking activity triggered by the ECG R-peaks, sorted by firing rate. Each row illustrates the spiking activity 700 to 1400 msec following an R-peak, with rows ordered by the firing rate in this interval. The y-axis depicts the indices of cardiac cycles in the ordering. For instance, index 50 indicates a cardiac cycle with the 50th smallest rate; index 300 denotes 300th smallest rate (which makes it one of the highest rates in baseline). The rates for both conditions are shown in the middle axes: black circles denote baseline, gray triangles denote IHE.

The bottom panel of Figure 2 shows the distributions of the average firing rates over heartbeats. Note that the distribution is either unimodal or continuously decreasing, with the majority of the mass concentrated around the lower rates. All subjects exhibit skewed firing rate distributions, indicating a large number of cardiac cycles with region rates that are much higher than the distribution mean. The tails of the IHE distribution were heavier than those for baseline for all subjects, suggesting an increase in the quantity of the high-rate heartbeats during sympathoexcitation.

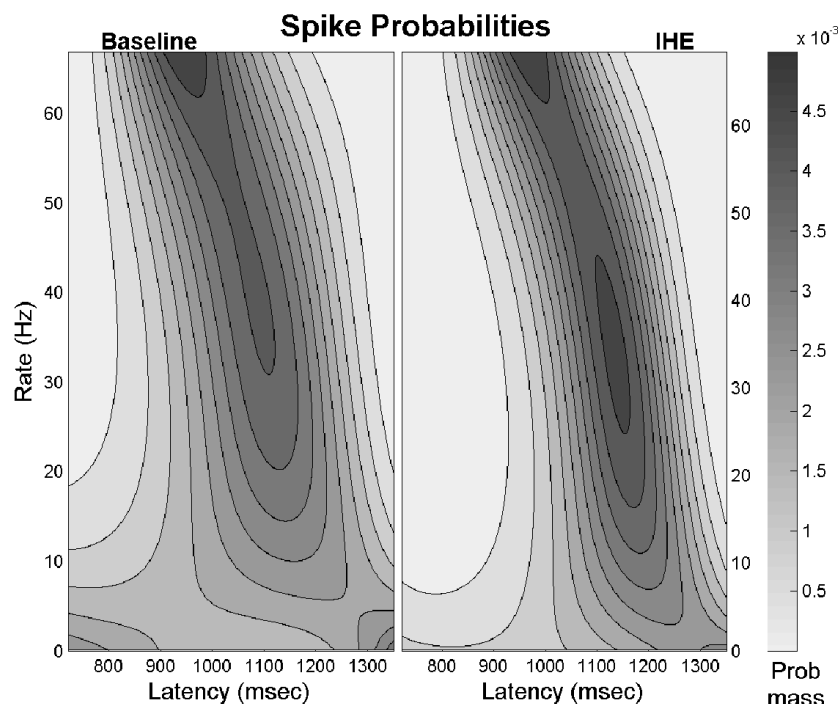


Figure 4. The fit to a model of the spike latency probability distribution as a function of the firing rate under baseline (left) and IHE (right) conditions. Similarly to Fig. 3, x-axis represents the latency within the high-density spiking region for each condition and y-axis represents the average firing rate within that region. Distribution is close to uniform for the low-rate regions and becomes increasingly concentrated about a preferred latency as the firing rate increases. Also, the preferred latency occurs earlier and the distribution becomes progressively more concentrated with increasing firing rate.

Overall, the visualization analysis revealed a similarly high variability in both baseline and sympathoexcitatory spiking. Our primary focus was the high-frequency within-heartbeat firing patterns, which varied similarly as a function of rates during the two conditions. Because of such high-frequency focus and because it was reflected in the average spiking rates, we did not explicitly measure condition-dependent low-frequency variability. It was, however, clearly present as there was an abundance of very low and very high rates of firing as well as a poor predictability of their occurrence in both conditions. Progressive sympathoexcitation was accompanied by an increase in the number of high-rate sympathetic bursts as well as by a rise in the firing rates within them.

These observations suggest that the sympathetic outflow intensity is governed by two factors: the adjustment of the firing rate within high-frequency bursts of spikes and the modification of the number of cardiac cycles containing such bursts relative to the number exhibiting low-rate firing or no activity at all.

2.3.2 Modeling

Extracting individual spikes granted us a closer look at the structure of the sympathetic bursts. Figure 3 highlights, for a representative subject's baseline (left, black) and IHE (right, gray) conditions, how the precise timing of spikes within a region of high frequency firing depends on the region's average rate. Specifically, the top panel displays the pronounced differences in the organization of the spikes between regions with low average firing rates and those with high average firing rates. The firing in the low-rate intervals is approximately uniformly distributed along the latencies, but spike distribution in the higher-rate regions acquires a Gaussian shape, with spikes concentrated about a particular latency. We constructed rasters sorted by the average rate within these intervals to evaluate the progression from the uniform to the concentrated firing states (see bottom panel of Figure 3). Specifically, as we move from the low to high-rate regions, there is a gradual emergence of spike concentration about a specific *peak latency*, which shifts closer to the corresponding R-peak as the average region rate increases. Importantly, the additional spikes that boost the average spiking rate occur predominantly around the peak latency, progressively increasing the spike concentration around it. A few spikes, however, occur further away from the peak latency, making the spiking bursts appear wider as a function of the increasing rate. Notice that this

progressive change in the firing structure of high-density spiking regions was not qualitatively altered by sympathoexcitation and was present during both baseline and IHE experiments for all subjects.

Subject Number	Baseline		IHE	
	KS Statistic	Critical Value	KS Statistic	Critical Value
1	0.0244*	0.0306	0.0359	0.0260
2	0.0312*	0.0323	0.0176*	0.0218
3	0.0270	0.0238	0.0268	0.0178
4	0.0196*	0.0419	0.0375	0.0232
5	0.0263*	0.0324	0.0278*	0.0319
6	0.0414	0.0297	0.0284*	0.0326
7	0.0439	0.0228	0.0277	0.0197
8	0.0456	0.0346	0.0302	0.0216

Table 3. Kolmogorov-Smirnov statistics and their corresponding critical values. For a model to pass the KS test, the statistic has to be smaller than the critical value. Models for 7 datasets passed the KS test (marked by asterisks). Note that the values of KS statistics are close to the critical values, indicating that deviations between modeled and actual spike placement probability were relatively mild.

We fit the model in Eq. 1 to the extracted spike trains using maximum likelihood; Fig. 4 displays the results for a representative subject under baseline (left) and IHE (right) conditions. Specifically, each row represents a region of intense firing with an average rate indicated on the y-axis. The colors indicate a probability density that a randomly selected spike occurs at a specified latency. The model captured two main spiking features. First, the structure of the low- and high-rate regions of high-density spiking differed. Specifically, the low-rate distribution had high variance that spanned most of the modeled region (see the light gray color spanning most of the bottom rate regions in Figure 4). In contrast, the spike location distributions for the high-rate spiking bursts were narrow; most of their mass was located around peak latency, indicated by the relatively narrow areas of dark gray near the top of Figure 4. Secondly, the transition

from low- to high-rate firing mode was achieved by a gradual constriction of the density and a progressive heightening of the probability distribution (see progressive narrowing of the dark gray area as rates increase; they also become darker for higher rates). Moreover, the location of the peak probability has progressively shifted to the shorter latencies.

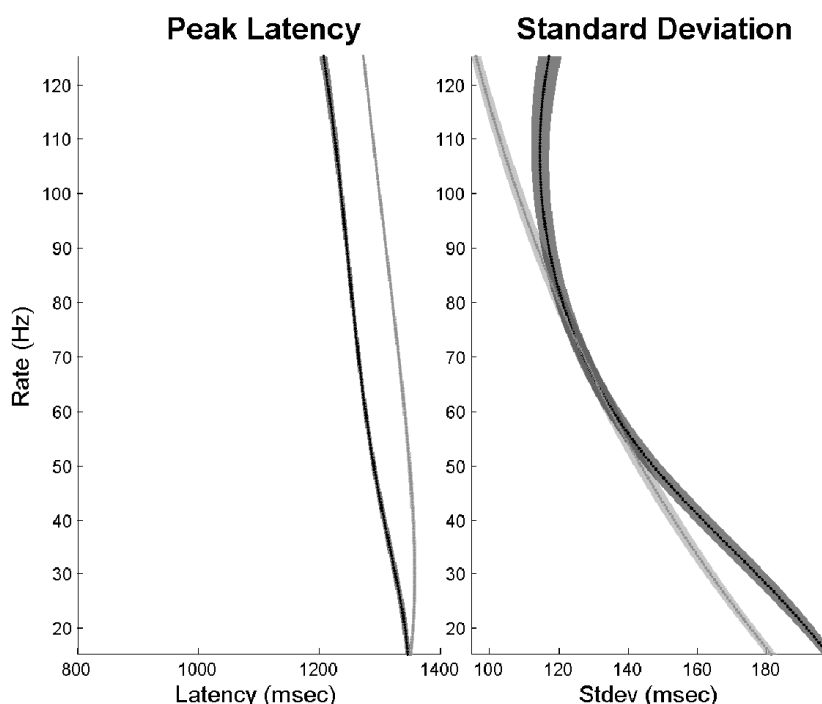


Figure 5. Features of the cardiac cycle spike latency distribution as a function of the firing rate. Subjects were pooled within baseline (black) and IHE (gray) conditions. The peak latency identifies the time after an R-peak with the highest intensity of spiking. In both conditions, peak latency decreased with the increasing firing rate; it was, however, significantly higher during IHE than at baseline for all rates. The standard deviation indicates the concentration of spiking about the peak latency. It is statistically indistinguishable between baseline and IHE for a large number of rates, but its reduction implies that spiking becomes more concentrated at higher rates in both conditions. Shaded regions represent bootstrapped 95% confidence intervals of the mean parameter estimates.

The probability structure was described using two features of the distribution: peak latency and standard deviation. The values of these features pooled for all subjects in the study within baseline (black) and IHE (gray) conditions are shown in Figure 5. For

both conditions, the peak latency shifts closer to the corresponding R-peak as the firing rate increases, indicating that the spiking starts earlier for the cardiac cycles with increased activity. The standard deviation of the spike locations becomes smaller with higher rates, indicating that the spikes are progressively more concentrated about the peak latency.

It is important to note that the two parameters used to describe the intra-region spiking vary identically as a function of the increasing firing rate for all subjects, specifically, the latencies get progressively shorter, and the standard deviation decreases. What is more remarkable is that this pattern is retained for two different states of the cardiovascular system: a dynamic equilibrium at baseline and a vigorous response to a strong sympathoexcitatory stimulus during IHE.

The expected latency showed significant differences between conditions. It occurred consistently earlier for baseline than for handgrip (see left panel in Fig. 5), indicating that the point of maximum spiking is attained later during sympathoexcitation than it is at rest. This is true for cardiac cycles at all firing rates. The standard deviation was statistically different for most firing rates, but the differences were rate-dependent, pointing to the lack of clear differentiation in spike concentration between the conditions.

The models' goodness-of-fit is described in Table 3. The left column for each condition contains the values of the KS statistic calculated from the data. Right-hand-side columns contain corresponding critical values: a maximal KS statistic value for the model to pass the test. Two things are worth noting. First, models for seven out of 16 datasets passed the KS test as indicated by their KS statistics being smaller than critical values.

The small differences between the KS statistic and the critical values suggest that the model fits for other subjects/conditions came close to passing the KS test. Second, the number of samples for each test is large (varies from about 600 to over 5000), which substantially shrinks KS statistic confidence intervals. This indicates that even a relatively mild deviation of the modeled from the actual probability distribution will force the KS statistic to become insignificant. While higher order polynomial models are likely to improve this goodness-of-fit marginally, no low-order polynomial model will provide a perfect fit to the data, and we expect KS tests to fail as the amount of data becomes large. Therefore, we opted to use a cubic model that parsimoniously described the main important features of the distribution, and which reliably and consistently produced parameter estimates with well-defined error bounds for all data.

2.4 DISCUSSION

The primary purpose of this article is to highlight the utility of using multifiber spike trains for assessment of sympathetic outflow to the vasculature. We showed that spiking is capable of expressing the aggregate features of sympathetic activity as well as traditional analysis methods are. Moreover, spikes provide information about sympathetic activity at a high temporal resolution. We used this information to assess the structure of the pulse synchronous bursts of spiking in terms of a peak systemic response and spiking concentration. In comparing statistical features of spiking activity at rest and during sympathoexcitation, we identified a number of striking similarities as well as one key difference between these conditions. Specifically, the latency of the peak response shortened and the spiking became more concentrated with increasing burst average rate

during both conditions. The peak response latency, however, was consistently and significantly longer during the sympathoexcitatory maneuver than at rest. These findings support the conclusion that similar mechanisms underlie the production of sympathetic spiking at rest and during excitation, and that these states differ only in the spike rates required to produce a particular spiking pattern.

2.4.1 Why Spikes?

The coarse features of the sympathetic outflow that have previously been detected using traditional analysis methods were evident in the extracted spiking. First, several researchers showed that sympathetic outflow progressively increases during a sympathoexcitatory maneuver [72, 106, 117]. In fact, both the number and the amplitude of the sympathetic bursts were shown to increase. Similarly, we see a progressive rise in the average spiking rates and the number of sympathetic bursts exhibiting high rates (>10 Hz) of firing. Second, we find the average latency of the sympathetic spike bursts to be close to the burst peak latencies detected in earlier studies: 1.27 (S.E.=0.05) sec for baseline and 1.28 (S.E.=0.04) for IHE, compared to 1.30 (S.D.=0.1, no S.E.) reported by Fagius et al. [41] and 1.33 (S.E.=0.03) sec reported by Wallin et al. [130]. Third, we see the peak latency decrease as a function of the average firing rate within the MSNA burst. Similarly, Wallin et al. [130] reported negative correlation between relative burst amplitude and burst peak latency.

Also, the variability present in the spiking was similar to the cardiovascular variability reported using traditional MSNA analysis techniques. For instance, we found a strong phase locking between sympathetic spiking and the cardiac cycle (see Figure 1).

Similar rhythmicity has been reported in both animals [1, 14] and humans [25, 26] using traditional analysis methods; in fact, cardiac rhythmicity was one of the earliest MSNA features detected, and it is used today to differentiate muscle from the skin sympathetic activity [85]. Also, variability of spiking across heartbeats was shown to have a slow rhythmic structure (with a period of approximately 10-15 seconds in Figure 1) consistent with frequency bands with increased power (0.1 and 0.25 Hz) reported by other researchers [113, 84].

In addition to the gross sympathetic outflow features, the extraction of individual spikes from the raw voltage trace granted us access to the information not previously available. Specifically, traditional methods involve a manual identification of the bursts of sympathetic activity based on their shape and amplitude; the statistics of these bursts are subsequently analyzed, while MSNA regions that fail the selection criteria are ignored. Here, we used the entire spike train. For instance, intervals that were thought to be completely silent due to systolic inhibition turn out to have sparse spiking (note spikes present at 400-800 and 1300-1700 msec latencies in Figure 1). In addition, we actively use the heartbeats that were traditionally dismissed as having too low or no sympathetic activity. Juxtaposing these areas with periods of high activity was key to detecting the striking similarities in how the fine temporal MSNA patterns change at rest and during sympathoexcitation. Specifically, we found that the latencies for the low-rate bursts are distributed uniformly: the spikes are equally likely to occur at any latency within a burst. As the firing rate increases, the spikes converge around a specific latency, and the spike latency distribution becomes more concentrated. Such differences in the distribution

structure suggest the presence of two classes of fibers: one that is active tonically at low rates and displays lesser strict cardiac rhythmicity, and another – reflexively activated – more closely phase locked to the cardiac cycle. This suggests in humans what has been shown in animals: that cardiac rhythmicity might be weak in up to one fifth of sympathetic neurons [51, 64]. Additional studies (e.g. ganglionic blockade) would be needed to validate this hypothesis and to ensure that the spikes we observed are produced by the muscle sympathetic efferents and not by other fiber types (e.g. skin afferents) contaminating the recording.

The additional information we extracted from the spike trains yielded a much more precise and physiologically interpretable description of sympathetic outflow. The main improvement is the removal of the chronological distortions introduced into the MSNA signal by the temporal smoothing used in traditional analyses [53]. For instance, the conduction delay is usually defined as a temporal difference between the ECG R-peak and the top of its corresponding integrated neurogram burst [40, 25, 41]. The neurogram peak denotes a point where the MSNA starts to ebb just enough to offset the increase in the integrated neurogram signal due to sympathetic activation. This point can fluctuate from burst to burst and can be hard to interpret physiologically. In contrast, we measure the latency between the R-peak and the point when the spiking distribution is maximal. This makes the latency more physiologically interpretable: it is a delay between the initiation of a hemodynamic event from myocardial depolarization to maximum sympathetic response. This difference in the quantities measured likely also contributed to numerical discrepancies found between the latencies in our study and in the literature.

For instance, the intra-recording latency variability obtained with spiking is smaller than the one seen in a traditional MSNA signal. Specifically, the cross-subject mean latency range obtained with the spiking model is 79 (S.E.=20) msec for baseline and 86 (S.E.=19) msec for IHE. In contrast, Wallin et al. [130] reported a cross-subject mean burst peak latency ranges (obtained with linear regression) to be 200 (S.E.=20) msec.

2.4.2 Mechanisms and Significance

The similarities in spiking patterns between rest and sympathoexcitation were present in both the coarse features of the spiking and in the structure of the individual bursts. For instance, a large number of cardiac cycles with little or no activity occur during both conditions despite a substantial increase in the nerve outflow during sympathoexcitation; the burst rate variability remains large as well. The spiking structure of individual bursts is also similar: the mean of the spike latency distribution shortens and its variance decreases with increasing average rate. These many similarities suggest that the mechanisms producing these spiking patterns are similar.

Studies show that the primary source of variability in the sizes of the traditional MSNA bursts depends on the specific sympathoexcitatory reflex activated. For instance, burst size was increased by fiber recruitment during end-inspiratory apnea [122, 12], but not during moderate orthostatic stimulus [110]. Our results suggest that the increase in MSNA outflow during IHE (when compared to baseline) is triggered primarily by the increase in the output of the already active neurons. We hypothesize this to be the case because the shift of the rate/latency curve in the left panel of Figure 5 denotes that the rate produced by the fibers firing at a specific latency increases. The additional spikes can

be produced by either recruiting more neurons with the same latency or by increasing the output of the fibers that are already active. The former seems unlikely because fiber recruitment is suggested to occur in order from the slowest and smallest fibers to the fastest and largest [122, 110]. There is, however, strong evidence that various sympathoexcitatory maneuvers produce increases in the individual fiber output [83, 82, 129, 58], suggesting it to be the primary cause of the rate/latency curve shift. Murai et al. [92] came to a similar conclusion by analyzing the firing properties of individual fibers recorded with increased-impedance microneurography.

Our results also allow us to speculate about the mechanisms of the firing rate increases within a single experimental condition. Figure 3 shows the dense spiking to begin at earlier latencies for higher rates than for lower rates; the shifting peak spiking latency in the left panel of Figure 5 shows that our model captures these spiking dynamics. We also observe a decrease in the variance of spike probability p (Eq. 1) with rising rates, which denotes a progressively tighter organization of the spikes around the point of maximal spiking. Taken together, these observations indicate that as rates increase, a progressively higher proportion of the spikes occur at earlier latencies, and, therefore, the rate is controlled by these shorter-latency spikes.

There is conflicting evidence in the literature about the potential source for these spikes. On the one hand, some studies indicate that the neurons at shorter latencies are likely to be different from the neurons at higher latencies. Specifically, the size of the traditional MSNA bursts was found to correlate strongly with the number of action potential types present in the bursts [122, 110], increasing the probability that the low-

frequency spikes at high rates came from a distinct population of neurons. More importantly, all the studies that classified human sympathetic efferents and tested them for size/latency relationship found that larger action potentials tend to have shorter latencies [122, 110]. Because the size of the action potential correlates with neuron's axon diameter [23] the population of neurons contributing the most to high rates in our study might be larger than the neurons most active at the lower rates.

On the other hand, Macefield et al. [82] found large variability in the spike latencies for individual fibers (358, S.E.=33 msec), raising a possibility that short-latency spikes we observed at high rates came from the same neurons that produced spikes at longer latencies during lower rates. In the same study he found that the spiking onset latency for individual fibers correlates well with the onset latency of the traditional MSNA bursts. Since large bursts have short onset latencies and because high region rates in our study correspond to larger bursts in the traditional MSNA analysis, this suggests that the additional earlier spikes we observe causing rate increases might come from the fibers active in smaller bursts and not from the newly recruited ones. In the subsequent study Macefield et al. [83] showed that the MSNA increase seen in subjects with high spontaneous baseline activity vs. those with low such activity can be attributed to an increased number of heart beats that fibers contribute a single spike to. This provides further evidence that fiber recruitment is not responsible for MSNA increases within a single experimental condition.

Considering the conflicting evidence in the literature and because we did not sort the spikes, we hypothesize that both fiber recruitment and increased active fiber output

contribute to the rate increase we observe within individual experimental conditions.

We also hypothesize that the latency shift during sympathoexcitation has functional significance. Specifically, there exists a negative correlation between rates and peak latency (left panel of Figure 5), which implies that a recording with a higher average spiking rate should have a shorter average peak latency. However, as we see in Table 2, the average spike latency during baseline is statistically indistinguishable from that during sympathoexcitation. We believe that by shifting the entire rate/latency curve to the higher latency dynamic range, the cardiovascular system is able to compensate for the higher average firing rate during sympathoexcitation and maintain a consistent range of mean latencies across conditions. Our hypothesis is consistent with effects shown to occur during other sympathoexcitatory maneuvers: Salmanpour et al. [110] showed an existence of a consistent latency shift across multiple outflow intensities during lower body negative pressure. The potential latency difference between his baseline and LBNP recordings (as indicated by the range of latencies recorded across individual fibers) was 75 msec, while the actual latency disparity (measured by average burst peak latency) was 10 msec. End-inspiratory apnea exhibited an analogous effect with a 100 msec potential [122] and 10 msec actual latency [40]. Similar situations might arise in other sympathoexcitatory conditions that have small mean latency differences with baseline: cold pressor (0 msec), propranolol injections (10 msec), and sodium nitroprusside injections (20 msec) [40].

Our work does not directly address the physiology of the rate/latency dynamic range resetting, thus we can only broadly speculate about its mechanisms. Figure 5 shows

that out of three distribution parameters only the rate is altered, which requires the spike distribution for a particular rate to shift but to not otherwise transform. This suggests that any partial changes in spike production that modify the distribution (e.g. recruitment of *some* slower fibers, neural pathways sharing, etc.) are not likely to cause the shift. A delay that is consistent for all rates is a more likely culprit. One such possibility is additional time required for central drive processing, especially at the end of handgrip when subjects must overcome fatigue. Switching of all processing to a slower central pathway (akin to the discrete latency variabilities during sinus nerve stimulation of vagal cardiac motoneurons [87, 88]) is another possibility. A study specifically targeting this issue is useful to clarify the mechanisms introducing this constant delay.

2.4.3 Conclusions

Overall, the use of the spike trains extracted directly from the raw multifiber voltage trace allowed us to assess the properties of the sympathetic outflow on both coarse and fine time scales. Access to individual spikes elucidated features of the sympathetic outflow that were previously inaccessible, such as the spiking patterns of the individual bursts and inter-burst spiking activity. It also revealed a presence of multiple categories of fibers, some of which are more explicitly locked to the cardiac cycle, and some of which are less so. Applying to spike trains a modeling method based on the theory of point processes provided new evidence about the mechanisms of human sympathetic control, specifically that the increase of spike production by the individual fibers contributes substantially to sympathoexcitation. These results represent a first

successful step in the utilization in the peripheral nervous system of analysis methods fruitfully used in analyzing the dynamics of central nervous system spikes.

3. RELATING MSNA TO HEMODYNAMIC VARIABLES

Muscle sympathetic nerve activity (MSNA) is traditionally studied with a multifiber signal called *integrated neurogram* that clearly identifies bursts of activity originated by the cardiac cycle hemodynamics. Because of temporal and cross-fiber smoothing, the heart rate forms a hard upper limit on the frequency content that integrated neurograms can yield. Identifying individual spikes within the multifiber voltage trace increases the temporal resolution beyond the heart rate limit and exposes the structure of individual pulse-synchronous bursts of sympathetic activity. In this chapter we use spikes to create a statistical model that captures both the slow variability in sympathetic output and its fast-changing characteristics. We relate the sympathetic spiking to concurrently collected hemodynamic variables across multiple experimental conditions and suggest physiological mechanisms that can contribute to the MSNA variability we observe.

The differences between the integrated neurogram and the individual spikes have shaped not only the research questions addressed by the techniques, but also the MSNA frequency bands the two methods primarily focus on. We review the integrated neurogram first. It is created by percutaneously inserting a stiff tungsten microelectrode to into a nerve bundle that contains muscle sympathetic efferent axons. Extracellular multifiber voltage traces are recorded by the electrode, rectified and convolved with a leaky integrator kernel (0.1 sec time constant) [53]. This processing technique produces prominent peaks representing the sympathetic response to individual heart beats' hemodynamic events; they occur at a fixed latency to the ECG R-peak of their initiating

cardiac cycles. In frequency domain these bursts of activity correspond to a substantial increase in total power over a frequency band centered at the subject's heart rate (about 1 Hz at rest); it is often the largest such power increase of the integrated neurogram's power spectral density. The use of leaky integrator renders the spectral content at frequencies larger than heart rate uninformative, but frequencies slower than heart rate are often studied. Two such frequency bands often feature increased power and have been studied in detail. The first, called the Mayer wave frequency (LF), is usually centered at about 0.1 Hz in humans. MSNA LF variability is tightly coupled to the Mayer waves in the arterial blood pressure, leading some researchers to suggest that they both are driven by a central oscillator [91, 90] and others to propose that Mayer waves are transient resonances within the baroreflex feedback loop [67]. The second band of interest is centered at the subject's respiratory frequency (HF, approximately 0.25 Hz in humans at rest). The increase in MSNA variability at the respiratory band is physiologically interesting because not only is it caused by the mechanical changes in lung volume during respiration [76], but it is also regulated by the chemical and metabolic demands of active muscle [27], making it useful in studying exercise.

Extracting individual spikes from the raw multifiber voltage trace opened up access to MSNA frequencies higher than the subject's heart rate, and the structure of the pulse-synchronous MSNA bursts has been the primary focus of studies that utilized spike trains. The most frequently asked question was the following: as sympathetic output increases, do individual human sympathetic fibers have a specific recruitment order? The answer seems to be "yes:" Steinback et al. [122] classified action potentials during

baseline and inspiratory apnea and showed that larger spikes are only present during larger bursts, implying that larger fibers activate only during high sympathetic output. Other groups used similar approach to suggest that a fiber recruitment order exists during other sympathoexcitatory maneuvers such as the lower body negative pressure [110], Valsalva maneuver [108], and voluntary breath hold [12].

In addition to describing high-frequency activity, spike trains also carry information about the slow oscillations in the sympathetic output. The LF and HF sympathetic variability has not been studied nearly as extensively using spikes as it has with the integrated neurogram. To our knowledge, there exist only two studies that used spikes for this purpose. The first one, by Brychta et al. [16], captured the fluctuation of the power spectral density of MSNA spike trains as a function of neck suction stimuli of different frequencies. The second one, also by Brychta and colleagues [17], constructed a model that explained the LF oscillations in systolic blood pressure in terms of the LF MSNA variability and systolic blood pressure HF oscillations in terms of the HF respiratory signal. Neither study used the spikes directly; instead the spike trains were convolved a Gaussian kernel and the resultant nonparametric rate was used in model construction and for statistical inference. While this approach grounds the power spectral density computation in physical units of $(\text{spikes/sec})^2/\text{Hz}$ (as opposed to the "arbitrary units" of the integrated histogram), the smoothing inherent in the spike rate computation eliminates access to the high-frequency information in the same way that the convolution did in the traditional setting.

In this chapter we assess the sympathetic spike train variability at different frequencies and explain it either as a function of time or in terms the concurrently collected hemodynamic variables. Statistically we aim to capture and predict as much variability in the spike trains as possible: the long-term time-dependent sympathoexcitation caused by isometric handgrip exercise (IHE) (and the absence thereof during baseline), the slower-than-heart-rate oscillations that are coherent with hemodynamic variability, and the spiking structure within individual heart beats. The scientific goal of the chapter is to differentiate and compare the size of the effects that covariates have on the sympathetic spiking: separate long-term temporal MSNA increase from the LF and HF fluctuations, identify the best predictors for spiking at each frequency band, and separate the effects of different covariates at identical frequency bands.

We expand on our earlier work in [135], where we first identified the high beat-to-beat MSNA variability, but chose not to model it, focusing instead on the interaction between the within-heartbeat spiking and the intensity of total output for each heartbeat. Here, we decompose the beat-to-beat variability onto different frequency bands and attempt to explain the spiking within each band by some combination of covariates: recording duration, diastolic blood pressure (DAP), respiratory signal, spikes' latency from the corresponding R-peak. We only seek to capture and characterize the spike train variability within each frequency band, therefore we neither model the interactions between spiking at different frequencies, nor do we look at the cross-band influence of the model covariates.

We structure Chapter 3 in the following manner. In the Methods section we describe how we record and process data and discuss our choice of the model parameters. We also describe our modeling methodology and specify the model formulation. In the Results section, we use visualizations to articulate why we chose specific relationships to model as well as elaborate on how the parameters we fit informed us about the properties of the sympathetic nerve output. We focus on the contribution of each frequency band to the total explanatory power of the model and elaborate on the relationship between the LF and HF spiking to the hemodynamic variables. Finally, the Discussion section outlines the significance of our work within what we know about human cardiovascular physiology and compares it with similar studies by other researchers.

3.1 METHODS

This section outlines methods that we used in the analysis. First, we describe the subjects and data collection protocols, followed by the explanation of the spike extraction procedures. Second, we detail signal processing techniques used to compute spectral information and visualize spike/covariate relationships. Third, we discuss the covariate extraction and spike train modeling methods involved in assessing the predictive power of hemodynamic variables across different frequency bands.

3.1.1 Subjects, Measurements, and Protocols

All volunteers gave their written informed consent to participate. All protocols were approved by either the ethical committee at the Grenoble University Hospital Center or institutional review board at Spaulding Rehabilitation Hospital, and conformed to the Declaration of Helsinki.

We studied healthy men (n=3) and women (n=4) with no signs or symptoms of hypertension, diabetes, cardiovascular or neurological disease or cancer, with normal resting ECG, and no regular use of tobacco. The population statistics are available in Table 1.

Electrocardiogram was recorded using a standard lead II ECG. The arterial blood pressure waveform in a finger of the hand was derived on a beat-to-beat basis by the Finapres Blood Pressure System (Finapres, Ohmeda), with brachial oscillometric blood pressure (Dash 2000, GE) employed as a calibration for the continuous Finapres measures. Respiration was a) collected by placing respiratory bellows around the chest to determine breathing depth and frequency or b) derived from inspired volumes recorded throughout all protocols via an infrared analyzer (VacuMed) connected to a mouthpiece with two-way respiratory valve.

Baseline activity data was derived from a five-minute period of supine rest. Sympathoexcitatory data was obtained by subjects performing isometric handgrip exercise at 35% maximal voluntary contraction (MVC) force sustained to fatigue. (The subject had previously performed three MVCs on a handgrip dynamometer to determine 35% MVC.) During the exercise, target force was displayed on a computer monitor. Subjects had continuous visual and auditory feedback from the investigators to ensure maintenance of target force until exhaustion, which was defined by the inability to maintain a target force for more than 3 sec despite continued effort and verbal encouragement [10].

Upon extraction and prior to applying the signal processing and modeling techniques, all hemodynamic data was downsampled from the original 10-40 kHz sampling rate to a 1 kHz sampling rate.

3.1.2 Spike Extraction

A sympathetic activity assessment technique was devised that discriminates neural firing from physiologically-irrelevant information, and which accurately identifies multifiber nerve activity independent of any a priori criteria [124]. This method is based on the statistical characteristics of the different artifacts, noise, and multifiber nerve activity that compose the unfiltered recording. A variant of the Expectation-Maximization algorithm was used at successive stages of processing to separate these components to extract the multifiber spike train. The details of the algorithm are described in [124], and, briefly, consist of the following steps: line noise and muscle twitches are removed from raw voltage recordings, producing a set of action potentials. These are then treated as a convolution of a set of pulses (Kronecker delta functions) with a prototypical spike template, which is extracted from the data. The posterior probability of a spike given the data is computed at each sample by deconvolving the spike template from the signal and removing the background noise using a mixture of distributions model. The posterior probabilities are then thresholded and binarized, producing a spike train at 40 kHz. Note that by using the action potential template we assumed that the shape of a subset of action potentials in the signal is roughly similar and that some fraction of the action potentials is considerably larger in peak amplitude than the noise source. Beyond that, both spike shape and amplitude are explicitly ignored because they

are heavily influenced by the distance from the axon to the electrode and are highly variable across recording sessions. The output of the model, therefore, is a binary multifiber spike train where neuronal firings by different axons cannot be differentiated. A refractory period was enforced by removing any of the spikes within 0.5 msec of each other. Average morphology templates were then collected and compared to noise samples and prototypical action potentials to ensure that spike locations in the binary train mark voltage deviations that are statistically significant and resemble actual spikes. We used this technique to extract multifiber spike trains from eight subjects under baseline and sympathoexcitatory conditions. Baseline activity was collected during supine rest, and an isometric handgrip exercise (IHE) served as a sympathoexcitatory stimulus. Once spikes were extracted, they were binned using 1 msec bin width in order for the sampling rate to match that of the hemodynamic variables.

3.1.3 Signal Processing

Diastolic arterial pressure (DAP) signal was obtained as follows: a) blood pressure waveform was low-pass filtered at 10 Hz to remove high-frequency fluctuations that included measurement noise, b) a minimum value of the filtered signal was computed for each cardiac cycle, and c) a waveform for each cardiac cycle was turned into a horizontal line with its ordinate at the minimum value. This produced a tachogram-like step DAP function with same duration and sampling frequency as the original blood pressure waveform.

Power spectral density (PSD) and magnitude-squared coherence were computed using the multitaper approach [125] with 19 tapers and a time-bandwidth product of 10.

Briefly, each spike train is multiplied by several *tapers* (discrete prolate spheroidal sequences that are orthogonal to each other [119]). Each product is then Fourier transformed and their average is taken to obtain a multitaper spectral estimate of the spike train. This procedure decreases the variance of the estimate at the expense of local frequency resolution, which depends on the dataset length. Applying 19 tapers produced frequency resolution of 0.02 Hz to 0.05 Hz, which struck a good balance between the visibility of salient spectral features and the lack of overwhelming details for a frequency band of lower than 0.6 Hz. We expressed PSD in decibels: $db = 10 \cdot \log_{10}(D/D_0)$, where D_0 is a PSD reference value, in our case the mean spiking rate for the whole recording. PSD for IHE was expressed using baseline D_0 so that PSD values could be compared directly across conditions.

3.1.4 Occupancy-normalized histogram

Extracting individual spikes allowed us to analyze the sympathetic activity using visualization techniques designed specifically for spiking data. Occupancy-normalized histogram is one such technique. For a concurrent spikes-signal pair it is constructed as follows. First, a covariate histogram is computed such that bin edges specify ranges of covariate values that fall within each bin. Second, because each spike time corresponds to a covariate value at that time, total number of spikes/bin can be computed by counting how many corresponding covariate values fell into each successive bin. Third, divide each bin's spike counts by total amount of covariate samples a bin contains to obtain bin's firing rate in spikes/sample. Convert rates to Hz.

3.1.5 Phase extraction

Several ways to model spiking at different frequency bands exist; we chose to use phase as a covariate because it allows us to focus on the variability of the spike train within the frequency bands without accounting for the changes in the signal amplitude. We used instantaneous phase to model LF and HF variability, and we used position within the cardiac cycle (*cardiac cycle phase*) to model variability at and above the heart rate.

The LF and HF frequency bands were identified for each subject and experimental condition in the following manner. First, a multitaper estimate of magnitude-squared coherence was obtained as described above. Second, areas of statistically significant coherence were selected within [0, 0.6] Hz frequency range; peaks of those high-coherence areas were visually identified. A Gaussian curve was fit to the shape of the high-coherence region using least squares methodology: frequency was treated as an independent variable, and curve's mean was constrained to lie at the region's peak. Frequencies that fell within one standard deviation of the mean were selected for phase extraction.

The instantaneous phase is a continuous, monotonically increasing time-varying signal that can be compared to spike rate via trigonometric sin and cos functions. Once the bandwidths of interest were identified, instantaneous phase was extracted from the original signal as in three steps. First, the real-valued signal was band-pass filtered at the frequency band of interest. To preserve the predictive nature of our spike train model, a causal filter was employed. Second, a Hilbert transform was applied to the real-valued

filtered signal to obtain an imaginary time series that together with the real-valued signal formed a complex analytic representation of the original time series. Third, the instantaneous phase was computed as an angle between the positive real axis and the complex vector at each time point.

To compute cardiac cycle phase we define a cardiac cycle to lie between two consecutive ECG R-peaks. Phase is computed by counting the consecutive indices of each individual sample within a cardiac cycle. We start counting at the first sample after the R-peak and end counting at the very top of the next R-peak, producing a set of natural numbers going consecutively from 1 to the total number of samples in the cardiac cycle. This set is then scaled to have a range of $(0, 2\pi]$ radians, so that both 0 and 2π fall on the two consecutive R-peaks. The counting and scaling procedure is then repeated for all cardiac cycles within the recording, producing a right-continuous sawtooth function with a $(0, 2\pi]$ range that describes the relative position of each sample within a cardiac cycle. Please note that this procedure only operates on full cardiac cycles; partial cycles at the very end of the recording are discarded.

3.1.6 Modeling Techniques

We constructed a probability model to capture as much of the variability present in the sympathetic spiking as possible. We modeled the MSNA spike train as an inhomogeneous Poisson process that depended on a) the recording time, b) the instantaneous phase of the diastolic pressure and the respiration signal, and c) the cardiac cycle phase. Model parameters were selected because they have been shown to correlate with sympathetic output, respectively, at time scales of the whole recording, at Mayer

wave and respiratory frequencies, as well as at the frequencies of heart rate and above. Specifically, it is well known that IHE induces a progressive time-dependent increase in the sympathetic output [117, 107]. The dependence of spikes on the cardiac cycle phase was shown by identifying a mean of a convex probability distribution in Chapter 2 and in [135]. It is also known that DAP and respiration affect sympathetic output and sympathetic spiking at frequencies slower than heart rate [59, 51]; their relationship is also depicted in Figures 7-8 in section 3.2.1.

Mathematical formulation of the model follows. Let $\varphi^R(t)$ be the respiratory and $\varphi^{DAP}(t)$ be the DAP instantaneous phase, let $\varphi^C(t)$ be cardiac cycle phase, let t be the latency relative to the associated R-peak, and let $\lambda(t)$ be the instantaneous firing rate function. We fit the following model of the $\lambda(t)$:

$$\lambda(t) = \exp \left\{ \begin{array}{l} \beta_0 + \beta_1 t + \beta_2 t^2 + \\ \kappa_{HF}^R \cos(\varphi_{HF}^R - \theta_{HF}^R) + \kappa_{LF}^{DAP} \cos(\varphi_{LF}^{DAP} - \theta_{LF}^{DAP}) + \kappa_{HF}^{DAP} \cos(\varphi_{HF}^{DAP} - \theta_{HF}^{DAP}) \\ \kappa^C \cos(\varphi^C(t) - \theta^C) \end{array} \right\}. \quad (2)$$

Here cardiac cycle phase $\varphi^C(t)$ and instantaneous phases $\varphi^R(t)$ and $\varphi^{DAP}(t)$ are time-dependent covariates and β s, κ s and θ s are parameters. Superscript R stands for respiration and superscript DAP indicates diastolic arterial pressure. Subscripts indicate frequency bands: HF is high and LF respiratory frequency. The exponential function guarantees that the spiking rate remains positive. Note that the dependence on respiratory LF instantaneous phase is absent because respiratory signal has little energy at the Mayer wave frequency band and because LF respiratory phase has little effect on the spiking rate.

The model parameters are fit by maximum likelihood, i.e. they are selected to maximize the probability of observing the data. For each covariate signal and frequency band combination θ s describe the phase of the maximal spiking and κ represent degree of spike rate modulation by the phase. Along with the maximum likelihood estimators for the model parameters, we compute their variance-covariance matrix. Using this information we obtain the analytical expression for the asymptotic joint probability density of the modulation strength κ and preferred phase θ for each signal/frequency band combination from which we then obtain 95% confidence intervals for κ and θ .

3.1.7 Statistics

Paired t-tests were used to determine statistical significance of phases and magnitudes of the parameters in model (2). Phases were shifted prior to testing to prevent wrapping affecting the results.

3.1.8 Model assessment

We assess the model fit in two ways. For a visual assessment we compare the spiking rates generated by the model to the true spiking rates of the original MSNA spike train. The rates are computed by convolving spike trains with a Gaussian kernel. MSNA spikes are used to produce the true rate. The model rate is computed in two steps: first spikes are generated from the model $\lambda(t)$ in (2) by thinning [28], and second, these spikes are convolved with the Gaussian kernel to produce the instantaneous rate. Two differently varying rates are produced by kernels with two different widths: 10% and 1%

of the length of the recording. Identical kernels are produced to compute both true and model rates.

To qualitatively assess the amount of variance in the MSNA spiking that our model captures, we used a generalization of the coefficient of determination metric R^2 that is appropriate for our exponential model (2). We formulate this concept using

$$\text{deviance: } D(N_{0:T}) = -2 \log \frac{\Pr(N_{0:T} | M)}{\Pr(N_{0:T} | M_f)} = 2 \sum_{i=1}^T \left[\log \frac{\Delta N_i}{\lambda_i \cdot \Delta} - (\Delta N_i - \lambda_i \cdot \Delta) \right],$$

where T is the number of time steps in the recording, $N_{0:T}$ denotes the entire spike train, \Pr stands for probability, M denotes a model to be tested, and M_f refers to a full model containing one parameter for every observation so that the data are fitted exactly. λ_i is a value of the instantaneous rate (2) at time step i , $\Delta = 1/T$ is the discretization bin size, ΔN_i is the number of spikes at time step i . The pseudo- R^2 metric is then $R_L^2 = \frac{D_0 - D_M}{D_0} \cdot 100\%$,

where D_0 is deviance for the null model M_0 , which only contains a single constant term [24]. R_L^2 in this context denotes the percentage of variance in the true spikes explained by the model. In our case M_0 is a homogeneous Poisson model $\lambda_0(t) = k_0$, and M is $\lambda(t)$ in (2).

We also use R_L^2 to separately assess the explanatory power of three covariate classes by specifying the following models:

$$\begin{aligned} M_t &= \exp\{\beta_0 + \beta_1 t + \beta_2 t^2\}, \\ M_I &= \exp\left\{\beta_0 + \kappa_{HF}^R \cos(\varphi_{HF}^R - \theta_{HF}^R) + \kappa_{LF}^{DAP} \cos(\varphi_{LF}^{DAP} - \theta_{LF}^{DAP}) + \kappa_{HF}^{DAP} \cos(\varphi_{HF}^{DAP} - \theta_{HF}^{DAP})\right\}, \\ M_C &= \exp\{\beta_0 + \kappa^C \cos(\varphi^C(t) - \theta^C)\}. \end{aligned}$$

Each covariate class here is intended to explain the variability of a specific frequency range: M_t captures variability on the scale of the whole recording (e.g. the progressive IHE-induced sympathoexcitation), M_l captures the slower-than-heartrate frequencies, and M_C assesses the fast within-cardiac-cycle dynamics.

3.2 RESULTS

Following the Box and Jenkins [11] model development paradigm, we start discussing our results with simple visualizations. They highlight spike train features and correlated them with the hemodynamic variables. We then explore in detail the magnitude and phase relationships between the spikes and the covariates that guided our model selection. Modeling results are shown last.

3.2.1 Visualizations

Frequency domain analysis provides information about the signal variability. Plotting the power spectral density (PSD) visually identifies frequencies where variability is high. Figure 6 displays a stereotypical PSD of the multifiber muscle sympathetic spike train for both baseline and isometric handgrip exercise (IHE) experimental conditions. Prominent baseline PSD increases are observed for several frequency bands: [0.07, 0.17], [0.22, 0.29], and [0.32, 0.42] Hz. The structure of the handgrip PSD is different; it has two major upswings centered at about 0.01 and 0.11 Hz followed by a slow decline for higher frequencies. The trajectory of the decline contains additional increases located at 0.23, 0.31, and 0.4 Hz.

Some regions of increased power match between baseline and IHE density traces. Specifically, within the regions with frequencies smaller than 0.3 Hz, there are two such

matches: one centered at about 0.11 Hz and another centered at about 0.25 Hz. They correspond to the frequency bands that were previously identified as physiologically relevant: Mayer wave frequency of about 0.1 Hz [67] and respiratory frequency, which is 0.23 Hz for the subject in Figure 6 and has a cross-subject mean of 0.25 (S.E.=0.015) Hz. These are the frequencies we will focus on in our analysis.

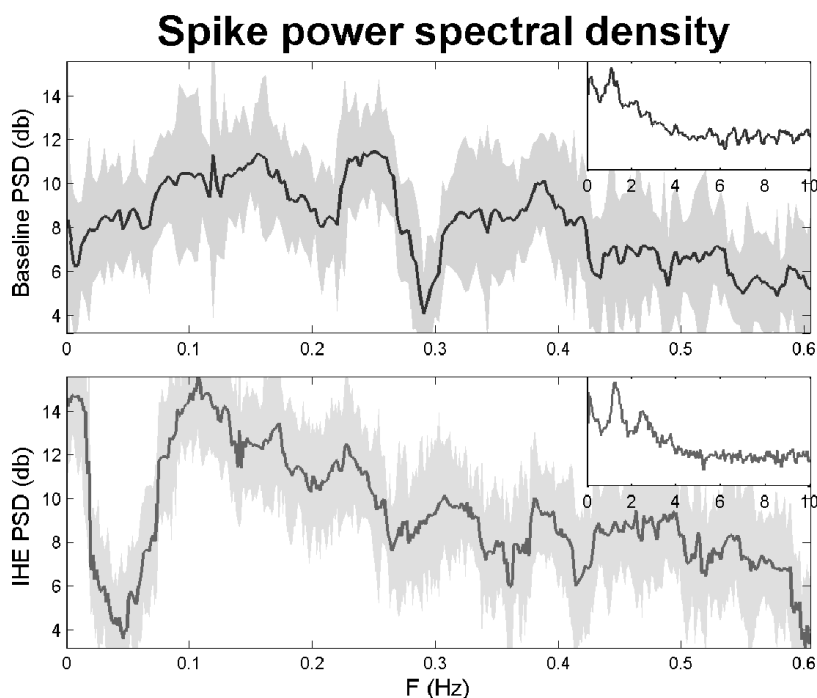


Figure 6. Multitaper estimates of the multifiber sympathetic spike train power spectral density (PSD) for a representative subject. Top panel (black) depicts baseline, bottom panel (gray) – isometric handgrip exercise (IHE). Shaded regions represent 95% confidence intervals. Top-right corner insets show zoomed-out PSD for frequencies of up to 10 Hz using units identical to those of the large plots. Baseline PSD increases at three frequency bands: [0.07, 0.17], [0.22, 0.29], and [0.32, 0.42] Hz. PSD for IHE is structured differently; it moves sharply up for a band of [0.08, 0.13] Hz and drops progressively down with upswings at around 0.23, 0.31, and 0.4 Hz. Note that the first two increases listed above for both experimental conditions correspond to the previously reported physiologically relevant frequency bands: Mayer wave (LF) and respiratory frequencies (HF). Also note that IHE exhibits a large increase in power at frequencies <0.02 Hz, likely corresponding to a progressive increase in sympathetic output induced by the isometric exercise. Insets show that PSD of the spike trains decreases for growing frequency bands, asymptoting to the mean firing rate. A sharp peak at the cardiac rate (around 1 Hz) for both baseline and IHE PSD indicates strong cardiac rhythmicity for both conditions.

The power spectral density as expressed in Figure 6 can be interpreted as showing the proportion of the total spike rate variability explained by the oscillations at a specific frequency. Frequency bands with high power explain more spike rate variability than do bands with low power. This interpretation suggests that the two sets of physiologically significant LF and HF sinusoids explain a large portion of the spike rate variability within the sympathetic spiking. We focus on assessing how hemodynamic variables influence the spike train oscillations at LF and HF in an effort to shed light on the possible physiological mechanisms that control slow variability in muscle sympathetic nerve spiking.

There are several ways to explore how covariates affect the spike train variability at different frequencies. A simple empirical approach is to construct a phase-based occupancy-normalized histogram: a bar plot of mean spiking rates for certain ranges of the covariate's phase. Figure 7 depicts such occupancy-normalized histograms for LF and HF bands for a representative subject. Each plot in Figure 7 is a snapshot of how much the spiking is varying as covariates' phases change.

All four plots in Figure 7 contain phases with higher spiking frequency when compared to frequencies of nearby phases; this effect is known as phase tuning, and the phase with the maximal rate is termed preferred phase. The strength of such tuning, exemplified by the heights of the peaks relative to the depth of the troughs, is different for different plots. For instance, low-frequency (LF) respiratory plot (top left) barely shows any tuning to the phase. The LF DAP tuning, to the contrary, is very strong, with

up to three-fold difference in spiking rates between the trough and the peak of the histogram.

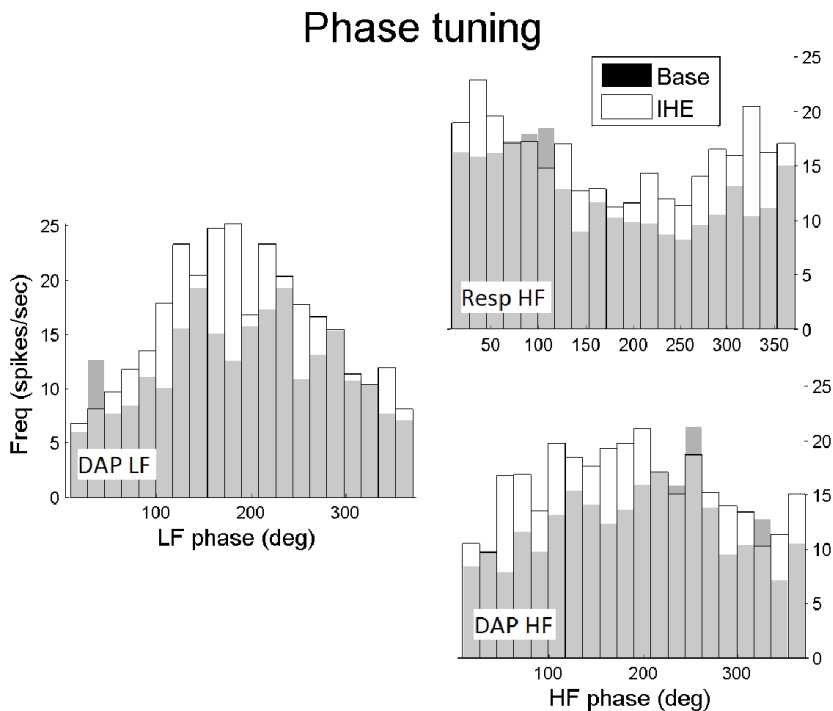


Figure 7. Frequency of sympathetic spiking as a function of covariates' phase for a representative subject. Left plot depicts spiking frequency dependence on low frequency (LF) phase of diastolic blood pressure (DAP). Top right panel displays spike rates as a function of diastolic arterial pressure (DAP) phase at respiratory frequency (HF); bottom right panel uses HF phase of respiratory signals. Filled bars represent baseline experimental condition; empty bars denote isometric handgrip exercise (IHE). Note that unshaded bars are generally taller for all phases, indicating sympathoexcitation during IHE experiments. Also note that sympathetic output is augmented by phases of all covariates; moreover, each signal exhibits a particular phase where spiking frequency is maximal (preferred phase). The most pronounced changes are observed in DAP LF plot (left), where preferred phases at both baseline and IHE occur at approximately 180 degrees, suggesting that the timing of DAP/spiking relationship at LF remains unchanged by IHE. Note also that preferred phases can change between conditions: the gray and white bars in both HF respiratory (right top) and HF DAP (right bottom) plots are similar in shape but are shifted relative to each other. This suggests that, while the spiking/covariate relationship remains quantitatively similar across conditions at HF, the timing of spike production changes with IHE-induced sympathoexcitation.

Juxtaposing shaded and empty bars shows how phase tuning changes from baseline to isometric handgrip exercise (IHE) respectively. Two types of change are evident: modifications in tuning strength or magnitude and variations in preferred phase.

For example, the relationship between the LF DAP phase and spiking (Figure 7, bottom left plot) does not change; for both baseline and IHE the sympathetic output is maximal at 180-degree phase, the point where DAP is lowest. However, the tuning magnitude is substantially stronger as evidenced by empty bars being taller than the shaded bars on that plot. HF relationships between spiking and both covariates are different: the strength of the relationship remains unchanged, but the preferred phase is substantially altered by IHE-induced sympathoexcitation.

In general, we see that hemodynamic covariates influence the spiking patterns within both low-frequency and respiratory frequency bands. Figure 7 shows that the relationship specifics depend on the frequencies and covariates in question. In addition, plots in Figure 7 indicate that sympathoexcitatory response to the handgrip exercise alters these relationships in a specific manner, only shifting the timing of some and only affecting the strength of the others.

Another way to assess the spikes-covariate relationship at different frequencies is coherence; it denotes a magnitude of linear relationship between the two signals at a specific frequency. Figure 8 displays values of spikes-DAP and spikes-respiration coherence for frequencies in the $[0, 0.6]$ Hz range.

The left plot shows the spikes/DAP coherence along with its 95% confidence interval. Coherence for both baseline and IHE is statistically significant in two frequency bands: $[0.05, 0.11]$ Hz (LF) and $[0.21, 0.28]$ Hz (HF). These bands also are nearly identical to the high-coherence bands that we automatically detected for this subject (shown in red): $[0.05, 0.13]$ Hz (LF) and $[0.2, 0.28]$ Hz (HF) respectively. The cross-

subject means of the automatically detected LF and HF coherent bands were 0.092 (S.E.=0.006) and 0.245 (S.E.=0.012). The centers of these bands lie close to the centers of the low- and respiratory-frequency bands that were previously reported to have physiological significance [96]. Note also that the spikes/DAP coherence values at LF are similar during both rest and sympathoexcitatory stimulation, but HF spikes/DAP coherence is higher at rest. This suggests that the linear relationship between the spikes and the diastolic pressure is not changed by the IHE-induced sympathoexcitation at LF, but it weakens at HF. The situation for spikes/respiratory coherence is similar, suggesting that sympathoexcitation might reduce the mutual effects that hemodynamic variables and sympathetic outflow have on each other at respiratory frequencies while leaving those effects unchanged for low-frequency bands.

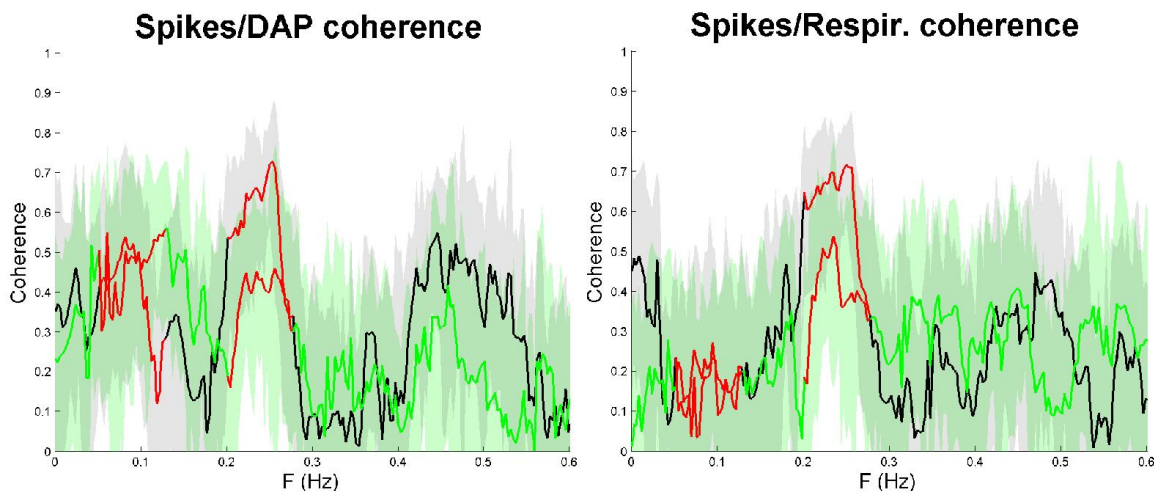


Figure 8. Multitaper estimates of magnitude-squared coherence between diastolic arterial pressure (DAP), respiration, and spikes. Left plot: coherence between multifiber sympathetic spike trains and diastolic arterial pressure for baseline (black) and IHE (green) conditions. Right plot: coherence between spike trains and respiratory signals for baseline (black) and IHE (green). Shaded regions represent 95% confidence intervals. Red indicates low-frequency (LF) and respiratory-frequency (HF) bands. Note statistically significant spikes/DAP coherence at LF (0.05 to 0.13 Hz) and HF (0.2 to 0.28 Hz) in both conditions. Significant spikes/respiratory coherence is only present at the HF

band in both conditions. In addition, both the spikes/DAP and spikes/respiration coherence is lower at HF during IHE, indicating the reduction of the association strength between spiking and the two signals at respiratory frequency during IHE.

3.2.2 Modeling Results

Here we introduce the results of modeling the sympathetic spiking using (2), which explains the variability of the MSNA spiking as a function of the covariates. We start with describing the dependence of the spiking on the instantaneous DAP and respiratory phase at low (LF) and respiratory (HF) frequency bands. For each instantaneous phase and frequency band, the results are described using two parameters: preferred phase, representing a part of the signal cycle that is predicted to have maximal spiking, and magnitude identifying the strength of the modulation. These parameters across the subject population are shown in Figures 4 and 5 as well as in Table 3.




Figure 9. Preferred phase model parameters for all covariates, subjects, and experimental conditions. Left-hand-side panel depicts low-frequency (LF) results; panels on the right show results at respiratory frequency (HF). Top panel shows respiratory signal preferred phase, bottom panels describe diastolic arterial pressure (DAP) preferred phase. Error bars are 95% confidence intervals; note that, because phase is circular, some confidence bounds wrap around either 0 or 360 degrees. Several things are worth noting. First, the consistency of LF DAP phases (bottom left) indicates that for all subjects the maximal spiking occurs at lowest values of DAP; this arrangement is also extremely consistent across subjects. Coupled with very small confidence intervals, this tight organization indicates that a large amount of LF MSNA variability is explained by the diastolic pressure. Second, the respiratory preferred phases for HF band (right top) also exhibit pronounced

organization: all (save two) the preferred phases fall within the interval [0, 150] degrees. Moreover, the preferred phases shift consistently later from baseline to IHE (the shift is statistically significant). Coupled with the lack of distinct organization and much larger confidence intervals of HF DAP preferred phases, the organization suggests that the majority of the MSNA variability in the HF band is explained by the variability in respiratory signal. Importantly, preferred phase shift indicates that while respiration still accounts for a large portion of the HF MSNA variability during IHE-induced sympathoexcitation, the MSNA/respiratory relationship changes with exercise onset.

Figure 9 displays the DAP and respiratory preferred phases. Bottom left plot is the most notable because of the consistency of LF DAP preferred phases for both baseline and IHE experimental conditions. For all subjects they are aligned at about 180 degrees (cross-subject mean 178.77, S.E.=7.07 degrees), which indicates a point of minimum DAP for each LF period. This is not surprising because a persistent negative correlation between MSNA and blood pressure was shown to exist during rest and other conditions [59, 105]. Surprising was the small degree of preferred phase variation across subjects. The standard deviation for baseline preferred phases was only 23.14 degrees, which using an average LF period of 8.7 seconds, converted to only 0.559 seconds! Moreover, handgrip did not change the precision of the tuning: cross-subject LF preferred phase standard deviation during IHE was mere 0.5861 seconds. Such small differences in MSNA relationship to LF DAP component across subjects and conditions is remarkable and suggest that there is a common underlying mechanism responsible for relating LF DAP to MSNA activation is extremely reliable and its function is unaffected by the sympathoexcitatory drive during the handgrip exercise.

Small confidence intervals for LF DAP phase (Figure 9 bottom left plot) indicate that tuning to phase is unambiguous and is likely strong. Indeed, Figure 10 shows that the tuning magnitudes for LF DAP are larger than those for HF DAP (LF DAP: 0.397

S.E.=0.295; HF DAP: 0.239, S.E.=0.508), indicating that DAP is the dominant covariate for the low-frequency and that it explains most of the LF spiking variability.

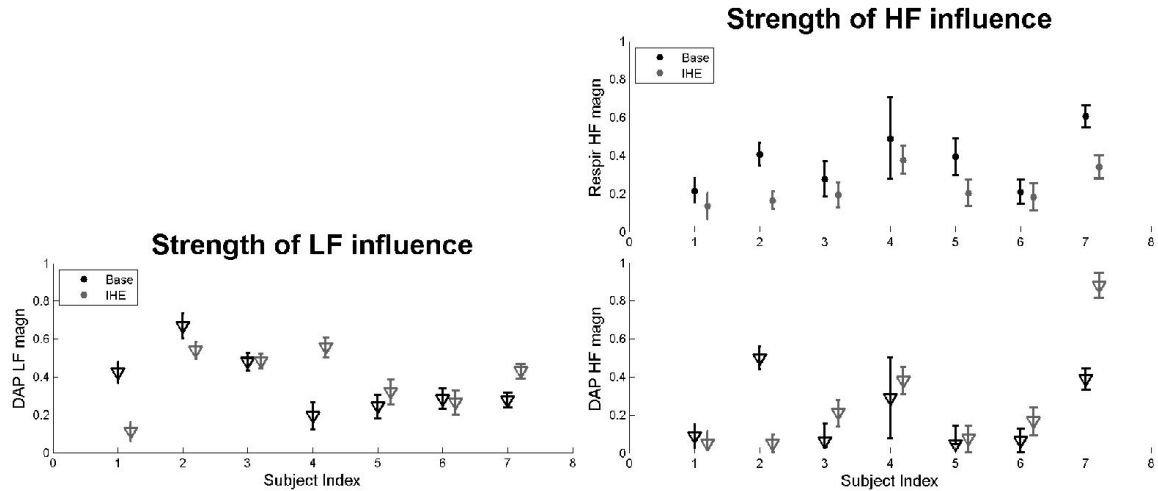


Figure 10. Modulation magnitude model parameter for all covariates, subjects, and experimental conditions. Left panel depicts low-frequency (LF) results; right panels show high-frequency (HF) values. Top panels describe respiratory signal magnitudes; bottom panels portray diastolic arterial pressure (DAP) magnitudes. Black markers indicate baseline values, and gray markers are isometric handgrip exercise (IHE) magnitudes; error bars identify 95% confidence intervals. This figure illustrates three interesting points. First, respiratory tuning magnitude at HF is higher for baseline than for IHE, but clear and consistent cross-conditional population changes are lacking for other covariate/frequency combinations. Second, DAP magnitudes are generally larger for LF than for HF independent of the experimental condition. Third, respiration is a dominant covariate for baseline condition at HF; dominance at HF during IHE is harder to identify.

Handgrip changes how much HF spiking variability each covariate explains. Specifically, it removes the dominance of respiration over DAP that is present at baseline and equates the explanatory power of the two covariates. This change is observed by comparing the respiration and DAP magnitude means: κ_{HF}^R is significantly larger than κ_{HF}^{DAP} at baseline, but they are statistically equal at IHE. The change can be attributed to the reduction of the respiratory tuning from Baseline to handgrip: mean κ_{HF}^R is significantly larger at Baseline as it is during IHE.

One of the simplest and most effective ways to show the very slow spike train variability is to convolve the spikes with a wide-bandwidth Gaussian kernel. The convolution produces a slow-varying trace that detects the global trends in the frequency of spike occurrence. Thick lines in Figure 11 depict this slow variability. They also show that our model is successful at capturing those long-term trends: the red thick lines run almost on top of the blue and black thick lines. Our model picks up both the stationarity of the baseline sympathetic output as well as the progressive time-dependent increase induced by subjects performing the isometric exercise.

MAGNITUDE PARAMETERS (unitless)						
	Baseline			IHE		
	LF	HF		LF	HF	
	DAP	Resp	DAP	DAP	Resp	DAP
Means	0.370	0.372	0.203	0.385	0.227	0.258
Stdevs	0.170	0.148	0.190	0.162	0.088	0.307
p-val vs. IHE	0.4312	0.0036	0.3092		0.3723	
p-val vs. DAP		0.0080				
p-val vs. HF	0.0408			0.1501		

Table 4. Cross-subject statistics for magnitude model coefficient κ across all conditions and frequency bands. p-values in shaded cells are <0.05 ; they quantify comparisons of the covariates' predictive power across experimental conditions, within and across individual frequency bands. First, magnitude does not change from one condition to another; the only exception is a decrease in high-frequency (HF) respiration magnitude during IHE (a lone shaded cell in "p-val vs. IHE" row). Second, $\kappa_{HF}^R > \kappa_{HF}^{DAP}$ at baseline, but the statistical significance is lost at IHE. Similarly, κ_{LF}^{DAP} is significantly larger than κ_{HF}^{DAP} at baseline, but no during IHE. Overall, these differences in tuning magnitudes suggest complex interactions between the spiking and the hemodynamic variables both within and across frequency bands.

By reducing the bandwidth of the kernel used for smoothing, additional variability can be picked up (see section 3.1.8 for kernel widths), as shown by thin lines in Figure 11, which even picked up the sympathetic bursts within individual cardiac cycles. More variable rates also expose our model's inability to capture some of the variability present in the spike trains. While accounting for large number of bursts that original rates

exhibit, the red thin traces fail to correctly capture the height of those bursts. This is unlikely to be an artifact of the kernel smoothing because identical kernel widths were used to produce rates from the original and modeled spikes. The more likely explanation is that the model cannot capture some proportion of the overall variance present in the sympathetic spiking. For instance, the sudden jumps in average burst rates highlighted by Figure 2 in Chapter 2 are likely unaccounted for by the phases of DAP and respiratory covariates we used in (2). Moreover, as we see in Chapter 2, the tuning of spikes to the cardiac cycle phase changes with time. We do not account for this change here, which can lead to improperly estimated burst heights.

Rates comparison




Figure 11. Comparison of nonparametric spiking rates computed from the original MSNA spike train and from spikes obtained from the model fit. Top plot shows baseline rates (black), bottom plot shows IHE rates (blue). Red denotes modeled rates on both plots. Each plot shows a rate produced

by a wide- and small-bandwidth kernel from the same spike trains. Note that the baseline wide-bandwidth rate remains steady, while the IHE wide-bandwidth rate climbs, reflecting the IHE-induced progressive increase in MSNA output. Also note that narrow-kernel rates capture substantial additional variability in the spiking, including individual cardiac cycle bursts. The model detects some of the higher-frequency variability, but fails to correctly pick up variation magnitude (red bursts are shorter than blue/black ones), indicating that some variability of the spike train remains unaccounted for.

We also quantified the proportion of the spike train variance that our model and its different covariate classes explain (see section 3.1.8 for methodological details). Figure 12 depicts the means of these percentages computed across the subject population. Notably, the cardiac cycle phase explains the most variance of all covariate classes, which is not surprising given the strong cardiac cycle dependence of MSNA. Recording time explains the smallest amount of variance, particularly for baseline, where spike trains are stationary. Note, however, that despite being small, the variance explained value for the baseline "Time" data point is statistically significant for 5 out of 7 subjects. The values for all other covariate classes and conditions are statistically significant save "Time" covariate class for subject 5. Also note that full model (2) explains on average less than 10% of the overall spike train variance. The large residual unexplained variance is a likely cause of the poor spike rate peak estimation in Figure 11.

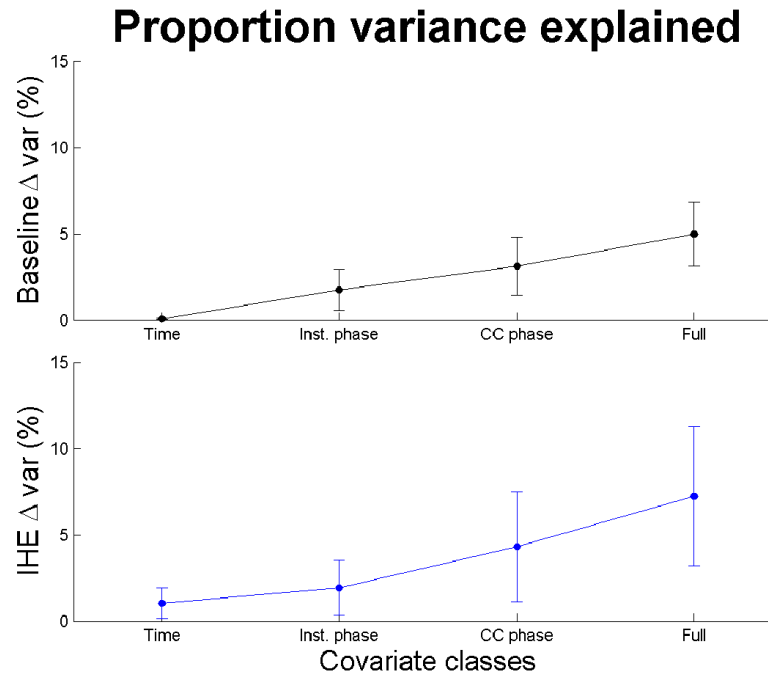


Figure 12. Proportion of spike train variance explained by different covariate classes from (2). Dots are means across subjects, whiskers are one standard deviation in length. Baseline (black) is shown in top plot, IHE (blue) on the bottom. "Time" represents time of the recording, "Inst. phase" denotes instantaneous phase of the DAP and respiratory signal at LF and HF bands. CC phase signifies cardiac cycle phase. Variance percentage is computed as in section 3.1.8. Note that recording time explains the smallest proportion of spike train variance, instantaneous phase covariates explain the next largest proportion, and the cardiac cycle dependence explains the largest proportion of spike train variance. The full model (2) explains under 10% of the total variance of the spike train. All deviance changes were statistically significant, except for two baseline and one IHE recordings. Deviance reduction used to compute the variance explained is significant for all subjects and covariate classes except the "Time" class for subjects 1 and 5 at baseline and subject 5 at IHE.

Overall, our results suggest that including many covariates from various frequency bands is a beneficial strategy in attempting to explain and model the very high variability present in the sympathetic spiking. For instance, our deviance analysis suggests that even when the dominant cardiac cycle effects are included, the others, operating on the slower time scales have considerable importance. However, those, operating on the time scale of the recording only gain substantial explanatory power if a stimulus directly effecting MSNA on that scale is applied. Within Mayer wave and respiratory frequency bands (the only ones we analyzed in detail), our data suggest that

each band has a dominant covariate that can predict a large part of the spike train variability within these bands. IHE perturbs this arrangement only slightly: LF tuning remains nearly identical, while the respiratory signal loses its predictive power with sympathoexcitation. The differential effect of IHE suggests that the physiological mechanisms responsible for associating the sympathetic output with the hemodynamic variables in humans might be different for different frequency bands.

3.3 DISCUSSION

In this Chapter we created a model that successfully captured some of the sympathetic spiking variability present in all physiologically relevant frequency bands. We described the slow trends in MSNA output that span the entire recording, we assessed the variations in spiking at the respiration and Mayer wave frequencies as well as at the frequencies of the heart rate and above. To our knowledge this is the first study that summarized the characteristics of the sympathetic output across that many frequency bands. The inclusion of the entire physiologically relevant frequency spectrum as well as the identification and explicit modeling of individual band's variability allowed us to disassociate the within- from the cross-frequency effects. Here we focused on within-frequency spiking/covariate relationships, identifying the effects several covariates have on the levels of MSNA output. Importantly, we disambiguated the effects of multiple covariates *within a single frequency band* and assessed how they differ with sympathoexcitation. In addition, we quantified the total proportion of variance present in the MSNA spike train that our model was able to capture.

Prior studies are inconclusive about how sympathoexcitation affects the relationship between MSNA output and the concurrently collected hemodynamic variables within individual frequency bands. The results vary based on which sympathoexcitatory maneuver is performed. For lower-body negative pressure (LBNP) Ryan et al. found that LF DAP/MSNA coherence increases during some pressures (45 and 80 mmHg), but not others [103]. Furlan et al. [43] and Kamiya et al. [73] showed that during heads-up tilt LF coherence between MSNA and arterial pressure consistently increases. Our results indicate that the spiking/covariate relationship is not only affected by the handgrip exercise, but the effect size changes depending on the frequency band of interest. First of all, we found that spike/DAP coherence was not affected by the IHE Mayer wave frequency (LF). Moreover, both the magnitude of the DAP phase influence on spiking and the DAP phase where maximal spiking occurs do not change. Therefore, our results strengthen the hypotheses that the effects of hemodynamic variables on the MSNA output are highly dependent on the exact stimulus used to produce sympathoexcitation.

The relationship between MSNA and covariates at respiratory frequency (HF) has also been shown to be experiment-dependent. For instance, both Furlan et al. [43] and Kamiya et al. [73] found that head-up tilt does not affect the values of HF coherence between blood pressure and sympathetic activity. Our results show that spiking/DAP coherence at HF decreases with IHE. However, because we used DAP and respiration signal as two separate covariates in our model, we can take this result one step further and disambiguate their relative contribution to HF variability within the spike trains.

Figure 10 and Table 4 show that IHE affects DAP and respiratory contribution to HF spiking differently. In particular, DAP magnitude parameter stays the same, indicating that DAP contribution is unaffected by IHE. The respiratory magnitude coefficient is diminished by IHE, indicating that respiratory signal loses some of its effects on spiking during exercise. These relative changes suggest a (at least partial) disassociation of the respiratory drive from the arterial pressure during the activation of the exercise stressor reflex.

Our work differs from the studies by Furlan et al. and Kamiya et al. in one more important aspect. In addition to using coherence, which is a frequency-specific descriptive method, we used a single model to capture the spike train variability at different frequencies. This allows us to compare the frequency specific model parameters and their relative changes across conditions. Table 4 shows one such change for the DAP magnitude parameter. In particular, the LF DAP magnitude parameter value is significantly higher than HF DAP magnitude at baseline. The statistical significance is lost during IHE. While the mean parameter values do not change that much from baseline to IHE (none of the differences are statistically significant), the loss of significance prompts further probing by, for example, explicit inclusion of the interaction parameters into the model.

It is well-known that lung volume is one of the primary causes of MSNA fluctuations at the respiratory frequency [27]. The origin of Mayer waves is not as clear. One theory states that LF MSNA variability is produced by a central nervous system oscillator that is independent of the baroreflex [67]. Under this theory, any changes that

IHE introduces to how baroreflex operates should not affect the MSNA activity at LF. Our results confirm this: neither the magnitude nor the preferred phase model parameters exhibit statistically significant changes at LF between baseline and IHE experimental conditions. The second theory is that LF variability represents a transient reverberation within the feedback loop of the baroreflex [67]. Under this theory, any changes to baroreflex that are introduced by IHE should be reflected in the strength (and potentially preferred phase) parameters of the model. We know that IHE dramatically changes the heart rate, altering the frequency of the cardiac cycle baroreflex feedback connections. This change would lead to altered reverberation frequencies and a decrease in the spiking/DAP association strength at LF, *which does not change from baseline to IHE*. We do not observe such a change, and, therefore, conclude that our results do not support the reverberation theory of Mayer wave generation.

3.3.1 Limitations

We use different respiratory signals for different subjects in our study. Five subjects use tidal flow, and two subjects use respiratory bellows signals. This places a potential limitation on our ability to interpret our results across the subject population. However, we believe that because of the covariates that we use in modeling and because of how we compare parameter values comparisons subsequent to the model fit, any adverse effects are minimized. Here is why. First, the phases of two signals correlate extremely well: signal peaks for both signals indicate top inhalation. Magnitude differences between the two signal types are irrelevant because we do not use that information in our modeling. Second, the top inhalation peak is flatter for the tidal flow,

which leads to a slight ambiguity in the identification of the precise maximal inhalation point. This ambiguity, in turn, can lead to a small shift in the instantaneous phase signal that we extract and use in our modeling. We do not believe that the phase shift impacts our analysis because we never directly compare HF preferred respiratory phase parameters across the population. We do compare magnitude parameter values, which would be unaffected by the instantaneous phase shift. So overall, we expect the usage of different respiratory signals to have negligible effect on the interpretability of our results across the population as long as we do not directly compare the respiratory phases extracted from the different respiratory source signals.

4 MARGINALIZED PARTICLE FILTERS

State space representation is a mathematical model of a physical system that uses latent state variables to describe how the system evolves in time and employs observation variables to represent the output of the system. In the statistical setting, both the state of the system and its output are random variables. The modeling is often done in discrete time by using difference equations to represent the system's temporal dynamics, with the system's output at each time step described by some function of state variables' values. If the state and observation equations are linear and the random variables are Gaussian, analytic solution is available in the form of a Kalman filter [69]. For more general systems, which are often nonlinear and non-Gaussian, analytic solutions do not exist. Differentiable state transition and observation functions that are mildly nonlinear can be linearized; Extended Kalman filter and its variants are based on such linearization [44]. For severe nonlinearities, sample-based approximation methods such as Gaussian Sum filter, Gaussian Quadrature filter, and Unscented Kalman filter are employed [2], [62], [68]. A major shortcoming of these models is a requirement that a probability density of the system state given the observations (*posterior* probability density function) remain unimodal [18]. More general techniques that deal with multimodal distributions were developed based on the Monte Carlo random sampling of the posterior [55, 54]. Particle filters (also termed Sequential Monte Carlo estimation) fall into this category of techniques. The main idea behind a particle filter (PF) is to create weighted random samples (*particles*) of the posterior distribution that accurately describe the state of the

dynamical system and to sequentially adjust them based on the incoming observations, all while maintaining the descriptive accuracy of the samples [46].

Numerous variations on the particle filter idea exist (see [18] for an excellent review). We will focus here on one specific modification that exploits an existing substructure within the dynamical system model; specifically, a division of the state random variables into a marginally nonlinear and a conditionally linear subsets. If such separation is possible, the nonlinear dimensions are marginalized out and estimated using standard particle filter techniques. The remaining linear dimensions are then conditioned on the nonlinear estimate, their noise is assumed to be Gaussian, and Kalman filter is applied to compute linear conditional estimate. This technique is called marginalization [116] or Rao-Blackwellization [21]. Marginalized particle filters (MPF) have been successfully applied to problems in aircraft [94] and underwater navigation [74], communications [22, 132], nonlinear system identification [79, 50], as well as audio source separation [5]. Some systems separate into linear/nonlinear components naturally. For instance, concurrent parameter and state estimation is often performed by adding parameters to the state vector, increasing its dimension. The estimation then proceeds as usual, but with the new, larger-dimensional, state vector. In order to allow reweighting and resampling to alter their values, the parameters are often modeled as a random walk with a small amount of Gaussian noise added at every time step [46, 77]. This is different from the usually nonlinear state transition functions, and such separation lends itself to the application of the marginalized particle filters, which have been extensively used in that setting [79, 80].

The separation of the posterior distribution into a marginal nonlinear and a conditional linear probability densities confers marginalized particle filters with several advantages over their standard counterparts. First, the overall variance of the particle filter estimator is reduced by the marginalization procedure. Specifically, when the PF estimate of the posterior density inference functions (e.g. mean, mode) is compared to its MPF estimate obtained by combining marginal nonlinear and conditional linear densities, the variance of the latter is smaller or equal to the variance of the former [33, 94]. Second, MPFs are likely to be more accurate than regular particle filters for the same number of particles. This is true because a) the nonlinear space necessarily has smaller number of dimensions than the total state space of the system, and b) the MPF particles live in the nonlinear dimension only, leading to its denser representation and higher quality inferences in the MPF case. Additionally, Kalman filter is the mean squared error optimal filter for linear Gaussian systems, so the estimation error within the linear dimensions is smaller for Kalman filters when compared to PF estimates. The third advantage is a performance increase that MPFs offer over some of their regular counterparts. These increases, however, need to be evaluated on a case-by-case basis because marginalizing as many states as possible will not necessarily produce performance gains. For instance, increasing the number of marginalized states for a large (7- to 9-dimensional) dynamical system produces only moderate gains in efficiency, and actually decreases performance at a certain point [75]. However, for both the smaller systems (3-6 total dimensions) and for the systems where observations do not explicitly depend on the linear states, increasing the number of marginalized states always led to

improved performance [75]. Notwithstanding the performance issues, marginalization remains for some large dimensional problems the only way to practically use particle-based estimation methods. For example, PF estimates of the integrated navigation system in the Swedish fighter aircraft Gripen, consisting of 27 total states, are too slow to be included in the on-board electronics. MPFs implemented in hardware, however, were successfully used because only three states are highly nonlinear [94, 116], which led to substantial performance improvement. Consequently, the variance reduction, more accurate estimation, and the performance gains make marginalization a highly successful addition to the sample-based estimation framework.

One common way to observe neural activity is to record the electrical signals produced by neurons and to extract individual spikes from the voltage trace. Because spike trains constitute discrete identical events, they can be modeled by point processes and used as observations in the state space estimation framework. Point processes necessitate highly nonlinear and non-Gaussian observation model, preventing the use of Kalman filters. Its analog for point process observations has been developed by Eden et al. [35] by using a Gaussian approximation to the posterior density. For multimodal posteriors a Gaussian approximation can lead to performance degradation, and methods that place no a priori assumptions on distributions should be used instead. Particle filters with point process observations (PFPPs) are an obvious choice; they have successfully been employed to estimate a variety of state processes from recorded spike train observations. For instance, the hand position, velocity, and acceleration traces were recovered from rhesus monkeys' premotor cortical activity [13] as well as from a

combination of premotor and somatosensory cortical recordings [134]. Cortical plasticity issues were addressed by using PFPPs to decode the temporal changes in the receptive field of hippocampal neurons from those neurons' spike train recordings [37]. In a novel application of PFPP, Meng et al. [89] estimated the static parameters of several biophysical models from the spiking data the models produced.

Real-life biological systems with point process observations present a challenge for sample-based dynamic estimation methods. The number of particles required for accurate estimation grows quickly with the increase in the system dimensionality. For instance, Wang et al. [133] used 100 particles to get as accurate PFPP estimates as she obtained with a state point process filter (SSPPF) [35] for a two-dimensional state. In order to get a 3-fold mean integrated squared error decrease from the optimal linear estimation, Brockwell et al. [13] had to use 2500 particles on a six-dimensional system. Wang et al. [134] used 1000 particles to improve SSPPF estimate by 15% for a seven-dimensional state vector. All of the examples above used simple zero-mean Gaussian random walk models to describe the evolution of the state variables through time. Realistic biophysical models are known to have highly nonlinear step-to-step transitions, which increase the complexity of the estimation problem. For instance, the Izhikevich neuron features a two-dimensional nonlinear dynamical system [63]. Neurons of FitzHugh-Nagumo type [93] are also small in dimensionality, but employ highly nonlinear transition functions. The well-known Hodgkin-Huxley system, detailing the cross-membrane currents that generate action potentials, consists of four coupled differential equations that account for membrane potential dynamics as well as time- and

voltage-dependent ion channels' gating variables. Using spikes to accurately estimate the state variables for such complex set of equations requires a very large number of particles, substantially increasing the computational burden. Indeed, for a full Hodgkin-Huxley+parameters system (a total of 6 dimensions) it took Meng et al. 10,000 particles to ensure that post-convergence voltage estimate lied entirely within the estimator's 95% confidence interval and that the particles for the concurrently estimated model parameters fully spanned their respective state spaces [89].

In this chapter, we aim to reduce the computational complexity and the number of particles required to produce accurate particle-based estimates of realistic biophysical systems. To achieve that goal, we expand the point process state space estimation framework to take advantage of the marginalization technique and define the marginalized particle filter with point process observations (MPFPP) algorithm. In the Methods section we describe the theory of particle-based estimation methods and of the marginalization procedure, as well as develop the marginalized particle filter methodology for models with point process observations. In the Results section, we apply the MPFPP estimation method to two models, a two-dimensional example with a switching nonlinear component and a full biophysical model of membrane potential. Restricting the first model to two dimensions allows us to easily visualize and study the probability densities arising during estimation. The abrupt switches between two nonlinear regimes also make it a hard decoding problem. The full biophysical model showcases some of the advantages of MPFPPs in terms of computational efficiency and estimation accuracy when compared with regular particle filtering. We then discuss the

implications of the marginalized filtering applied to point process observations in terms of variance reduction, estimation accuracy, and computational complexity topics discussed in the literature for the standard MPF. We will also mention the utility of the MPFPP method in tackling a problem of parameter estimation in complex biophysical models.

4.1 METHODS

We structure this section in a manner of a tutorial. We will first briefly cover the general theory of dynamic estimation and show where particle filtering fits within that theory. We will then describe how the particle filters can be augmented to take advantage of the linear substructure in the state transition model to produce marginalized particle filters for continuous state observations. We will then translate these ideas to point process observations and show that the theoretical leap between the two is rather simple.

4.1.1 Regular Particle Filter

The state-space estimation problem refers to methodology for approximating a continuous latent random process given a series of observations related to this process. In discrete time, the problem can be formulated as follows:

$$\begin{aligned}x_{k+1} &= f(x_k, \varepsilon_k^x) \\ y_k &= h(x_k, \varepsilon_k^y),\end{aligned}\tag{3}$$

where x_k is the state variable and y_k is the observation value in a discrete time bin k , ε_k^x is the process noise, ε_k^y is the observation noise, and functions f and h are arbitrary. The

noise densities $p(\varepsilon_k^x)$ and $p(\varepsilon_k^y)$ are independent and assumed to be known. The state and observation variables at each time step are distributed as $x_{k+1} \sim p(x_{k+1}|x_{0:k})$ and $y_k \sim p(y_k|x_k)$ respectively. We define the state trajectory to be a set of state variables $x_{0:k} = (x_0, x_1, \dots, x_k)$, and observation trajectory to be a set of observations $y_{1:k} = (y_1, y_2, \dots, y_k)$. x_0 is the initialization value for the states. The goal is to estimate the latent state trajectory $x_{1:k}$ from the observed trajectory $y_{1:k}$. This can be done by recursively computing the posterior density $p(x_k | y_{1:k})$ at each time step $k=1, \dots, K$:

$$\begin{aligned}
 p(x_k | y_{1:k}) &= \frac{p(y_k | x_k) \cdot p(x_k | y_{1:k-1})}{p(y_k | y_{1:k-1})} \\
 &\propto p(y_k | x_k) \cdot \int p(x_k | x_{k-1}) \cdot p(x_{k-1} | y_{1:k-1}) dx_{k-1}
 \end{aligned} \tag{4}$$

For the Gaussian noise terms and linear functions f and h , Kalman filter provides an analytic solution to the posterior density at each time step. No analytic solution, however, exists for the nonlinear functions of state. Several computational methods have been developed to deal with this issue; extended Kalman filter [65] and particle filter [46] are two of the most prominent ones.

Here we briefly describe the general principles behind the particle filter; various refinements are discussed in [18]. A particle filter (also known as Sequential Monte Carlo estimation) seeks to approximate the posterior distribution in (4) at each time step k with a set of values (called *particles*) $x_k^{(i)}$ sampled from that distribution. The particles are directly used to perform inference, calculate quantiles, etc. [31]. As the number of particles grows, the approximation becomes progressively more accurate. It is often not possible to sample from the posterior directly, so the particles can be sampled from

another distribution, called a *proposal* (or *importance*, or *instrumental*) density. To correct the discrepancy between the proposal and the posterior densities, a weight $w_k^{(i)}$ is associated with each particle $x_k^{(i)}$. The joint set $\{x_k^{(i)}, w_k^{(i)}\}_{1 \leq i \leq M}$ of particles and weights is then propagated from time step k to $k+1$ by passing the particles through a state transition function and updating each particle's weight to account for the likelihood of observing the data given the particles' state value. Further details are available in the Algorithm section below.

There are several ways to select the proposal distribution; here we use the state transition probability $p(x_k | x_{k-1})$. This selection makes our algorithm a bootstrap filter [46, 31]. It also simplifies the weight update, which becomes directly dependent on the observation likelihood:

$$w_{k+1}^{(i)} = w_k^{(i)} \cdot p(y_{1:k} | x_{0:k}^{(i)}) \quad (5)$$

Intuitively, the relative weight of each particle is increased or decreased dependent upon the relative size of the observation likelihood.

Due to their iterative nature, particle filters are susceptible to weight degeneracy. To deal with this issue, resampling was introduced by Gordon, Salmond, and Smith [46] to make the first operationally effective version of the particle filter. Here, we use a variant of the sampling importance resampling method from [101, 102] called *residual resampling* [80]. During residual resampling, the particles with large weight are reproduced, and particles with small weights either survive or die off with a certain probability; both the number of copies and survival probability depend on the particle's

weight. Mathematically, let M be the number of particles used. We retain $M_i = \lfloor Mw_k^{(i)} \rfloor$ copies of particle i , where $\lfloor \cdot \rfloor$ indicates rounding down to the nearest integer, and then obtain $M - \sum_i M_i$ i.i.d. draws from a set of particles $\{x_k^{(i)}; i=1, \dots, M\}$ with the probabilities proportional to $Mw_k^{(i)} - M_i$. After resampling, the weights of each particle are reset to $1/M$.

4.1.2 Marginalized Particle Filter

Frequently, systems in need of estimation are not entirely nonlinear and non-Gaussian; that is, there exists a subset of the state vector both linear and Gaussian conditioned on other states. Marginalization (also called Rao-Blackwellization) technique allows processing of such parts using standard optimal Gaussian filtering [116] while the remaining dimensions of the state vector are processed with a regular particle filter. We go over the basic principles of the marginalization technique below and then apply them to point process observations in the next section.

The main idea behind the marginalized particle filter is to take advantage of the conditional linear Gaussian substructure within the model and process the nonlinear and linear state subspaces separately. If a linear substructure exists, it can be marginalized out to obtain a marginal nonlinear posterior density as follows:

$$p(x_{0:k}^n | y_{1:k}) = \int p(x_{0:k}^n, x_{0:k}^l | y_{1:k}) dx_{0:k}^l.$$

This integral cannot be evaluated explicitly because we do not know the analytic expression for the joint posterior density. Instead, we use the regular particle filter to obtain estimates of the marginal nonlinear posterior; i.e. we restrict the set of particles and their corresponding weights $\{x_k^{n,(i)}, w_k^{(i)}\}_{1 \leq i \leq M}$ to operate in

nonlinear dimension only, allowing us to make inferences from this density's particle-based estimator. Once we obtain the marginal nonlinear estimate $\hat{x}_{0:k}^n$, we can use it to estimate the conditional linear posterior density as a random Gaussian mixture

$p(x_{0:k}^l | x_{0:k}^n, y_{1:k}) \approx \sum_{i=1}^M w_k^{(i)} p(x_{0:k}^l | x_{0:k}^{n,(i)}, y_{1:k})$, using the normalized nonlinear particle

weights $w_k^{(i)}$. Each particle's linear conditional density $p(x_k^l | x_{0:k}^{n,(i)}, y_{1:k})$ is computed

using standard Kalman filter recursions. The joint density then becomes the product of the nonlinear marginal posterior and the linear conditional posterior:

$$p(x_{0:k}^n, x_{0:k}^l | y_{1:k}) = p(x_{0:k}^l | x_{0:k}^n, y_{1:k}) \cdot p(x_{0:k}^n | y_{1:k}). \quad (6)$$

The joint density is seldom used, however, as the full state estimate can be reconstructed using the linear and nonlinear densities separately as follows:

$$\begin{aligned} \hat{x}_k^n &= \sum_{i=1}^M w_k^{(i)} x_k^{n,(i)}, \\ \hat{x}_{k|k}^l &= \sum_{i=1}^M w_k^{(i)} \hat{x}_{k|k}^{l,(i)}, \\ \hat{P}_{k|k} &= \sum_{i=1}^M w_k^{(i)} \left(\hat{P}_{k|k}^{(i)} + (\hat{x}_{k|k}^{l,(i)} - \hat{x}_{k|k}^l)(\hat{x}_{k|k}^{l,(i)} - \hat{x}_{k|k}^l)^T \right). \end{aligned} \quad (7)$$

$\hat{P}_{k|k}$ here is an estimate of the conditional linear posterior variance. Here, we expect the variance of estimators (7) to decrease as a function of increasing number of particles. We place $O(1/M)$ as a lower bound on the rate of decrease (see variance discussion in section 3.1 of [34] and in [33]).

The details of the marginalized particle filter estimation can be found in several publications, including [22], section IV of [32], as well as [6] and [116]. In the next

section we will develop the estimation details as they apply to filters with the point process observations.

4.1.3 Marginalized Particle Filter with Point Process Observations

Marginalized particle filters with point process observations (MPFPPs) differ conceptually very little from the regular MPFs: MPFPP observation density becomes Poisson to account for the point nature of observed events, a change that necessitates updates to the likelihood and linear estimation equations. Therefore, we will first define point processes mathematically so we can present the updated observation likelihood equation. We will then connect these mathematical descriptions to the filtering theory developed in the previous sections and provide a complete set of estimation equations and an algorithm, which together can be used to easily replicate the filtering procedure.

A point process is a type of random process describing a set of discrete isolated identical events, the timing of which is stochastic. Point processes are modeled by a parameterized instantaneous rate also known as conditional intensity function (CIF). Mathematically, let T be total observation time, during which J events occur at times $(0 < u_1 < u_2 < \dots < u_J < T]$. Also let $t \in (0, T]$, and let $N(t)$ describe a total number of events up to t . Allow $N_{0:t}$ to record precise event occurrence times $(0 < u_1 < \dots < u_j \leq t]$, and let it be called event history. Also define $\Theta = [\theta_1, \theta_2, \dots, \theta_p]$ to represent a set of P parameters describing the process generating the events, and x_t to denote covariates that event arrivals might depend on. Then the conditional intensity function is defined as

$$\lambda(t | x_t, \Theta_t, N_{0:t}) = \lim_{\Delta \rightarrow 0} \frac{\Pr((N_{t+\Delta} - N_t) = 1 | x_t, \Theta_t, N_{0:t})}{\Delta}, \quad (8)$$

where x_t denotes the values of covariates and Θ_t are values of parameter at time t . Note that defining CIF in this way grants us the flexibility of conditioning our instantaneous rate model on extraneous covariates, parameters, which might be dynamically changing, and event arrival history.

In order to model a point process, it is helpful to represent it in discrete time. To do so, we subdivide our time interval $(0, T]$ into K time bins each spanning a time interval $((k-1)\Delta, k\Delta]$ that contains ΔN_k events. We constrain ΔN_k to be at most 1 for all bins $k=1, \dots, K$ by selecting a very large K . We also define event history up to (but not including) bin k to be $N_{1:k}$, a discrete equivalent of $N_{0:t}$. In order to make notation less cumbersome, we define λ_k to denote the discrete version of covariate-, parameter-, and history-dependent $\lambda(t | x_t, \Theta_t, N_{0:t})$ in (8).

From the CIF definition above, we obtain an expression for probability of observing an event in a time bin k :

$$p(\Delta N_k | x_k, \Theta_k) = (\lambda_k \cdot \Delta)^{\Delta N_k} \cdot (1 - \lambda_k \cdot \Delta)^{1 - \Delta N_k} + o(\Delta),$$

where $o(\Delta)$ refers to higher-order terms with respect to Δ . For Δ small, this Bernoulli probability expression is well approximated as [15]:

$$p(\Delta N_k | x_k, \Theta_k) = (\lambda_k \cdot \Delta)^{\Delta N_k} \cdot e^{-\lambda_k \cdot \Delta}. \quad (9)$$

As we mentioned earlier, MPFPPs take advantage of the linear substructure within the dynamical system model in (1). For the most general case we can formalize the separation of the linear and nonlinear parts of the model as follows [116]:

$$\begin{aligned} x_{k+1}^n &= f_k^n(x_k^n) + A_k^n(x_k^n) \cdot x_k^l + G_k^n(x_k^n) \varepsilon_k^n \\ x_{k+1}^l &= f_k^l(x_k^n) + A_k^l(x_k^n) \cdot x_k^l + G_k^l(x_k^n) \varepsilon_k^l \end{aligned} \quad (10)$$

where x^n and x^l are linear and nonlinear components of the state vector respectively, functions f_k^n and f_k^l are nonlinear, A_k^n , A_k^l , G_k^n , and G_k^l are appropriately sized matrices that can depend on x^n , and the noise terms are Gaussian:

$$\begin{bmatrix} \varepsilon_k^n \\ \varepsilon_k^l \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} Q_k^n & Q_k^{\text{ln}} \\ (Q_k^{\text{ln}})^T & Q_k^l \end{bmatrix}\right), \quad \varepsilon_k^y \sim N(0, R_k).$$

The linear prior probability is also Gaussian: $x_0^l \sim N(\bar{x}_0, \bar{P}_0)$; the density of x_0^n is arbitrary, but assumed known.

The observations in our case are spikes, so observation random variables are Poisson-distributed (9):

$$\begin{aligned} y_k &\sim \text{Pois}(\lambda_k \cdot \Delta), \\ \lambda_k &= \exp\{h_k(x_k^n) + C_k(x_k^n) \cdot x_k^l\} \end{aligned} \quad (11)$$

where h_k is a nonlinear function, and C_k is an appropriately sized matrix.

As in the regular marginalized particle filter, the idea is to compute nonlinear marginal and the linear conditional posterior densities (6) separately for each time step. This is accomplished in two stages: prediction and update. Prediction step takes a guess at what the marginal nonlinear and conditional linear posteriors look like based on state transition functions only; the update step corrects this prediction based on the observed data. We'll focus on nonlinear particles first. To make a prediction from step k to step $k+1$, samples (nonlinear particles) from $p(x_{0:k}^n | y_{1:k})$ are passed through the nonlinear transition function (10), which produces samples of marginal nonlinear one-step

prediction density $p(x_{0:k+1}^n | y_{1:k})$. Particle weights are then updated to incorporate information about the observation at step $k+1$ by computing a marginal nonlinear likelihood of seeing a spike as in (9):

$$p(y_{k+1} | x_{0:k+1}^n) = p(\Delta N_{k+1} | x_{0:k+1}^n) = (\lambda_{k+1}^n \cdot \Delta)^{\Delta N_{k+1}} \cdot e^{-\lambda_{k+1}^n \cdot \Delta},$$

$$\lambda_{k+1}^n = \exp \left\{ h_{k+1}(x_{k+1}^n) + C_{k+1} \cdot x_{k+1}^l + \frac{1}{2} C_{k+1} \cdot P_{k+1|k} \cdot C_{k+1}^T \right\}. \quad (12)$$

This likelihood is a Poisson-distributed random variable with a new nonlinear rate λ_{k+1}^n that has been adjusted to account for likelihood's independence of the linear state dimensions. Updated weights make the correction step possible: the weighted samples now estimate the nonlinear marginal posterior $p(x_{0:k+1}^n | y_{1:k+1})$ and allow computation of point-mass estimates of the inference functions on that distribution. Alternatively, a shape of the marginal nonlinear posterior distribution can now be estimated/visualized using resampled particles. See sections II.A and II.B of [18] for a more thorough treatment of resampling.

Once the nonlinear estimates are obtained at either prediction or correction stages, the linear conditional densities for that step can be computed. For time step $k+1$, the conditional linear one-step prediction density $p(x_{k+1}^l | x_{0:k+1}^{n,(i)}, y_{1:k})$ for each particle $i=1, \dots, M$ is a Gaussian with mean $\hat{x}_{k+1|k}^{l,(i)}$ and covariance $P_{k+1|k}^{(i)}$. For a general system in (10) they are computed using Kalman filter-based recursion [116]:

$$\hat{x}_{k+1|k}^{l,(i)} = \bar{A}^{l,(i)} \hat{x}_{k|k}^{l,(i)} + G_k^{l,(i)} (Q_k^{\ln})^T (G_k^{n,(i)} Q_k^n)^{-1} z_k^{(i)} + f_k^{l,(i)} + L_k^{(i)} (z_k^{(i)} - A_k^{n,(i)} \hat{x}_{k|k}^{l,(i)}),$$

$$P_{k+1|k}^{(i)} = \bar{A}_k^{l,(i)} P_{k|k}^{(i)} (\bar{A}_k^{l,(i)})^T + G_k^{l,(i)} \bar{Q}^l (G_k^{l,(i)})^T - L_k^{(i)} N_k^{(i)} (L_k^{(i)})^T, \quad (13)$$

where

$$\begin{aligned}
N_k^{(i)} &= A_k^{n,(i)} P_{k|k}^{(i)} (A_k^{n,(i)})^T + G_k^{n,(i)} Q_k^n (G_k^{n,(i)})^T, \\
L_k^{(i)} &= \bar{A}_k^{l,(i)} P_{k|k}^{(i)} (A_k^{n,(i)})^T (N_k^{(i)})^{-1}, \\
z_k^{(i)} &= x_{k+1}^n - f_k^{n,(i)},
\end{aligned} \tag{14}$$

and

$$\begin{aligned}
\bar{A}_k^{l,(i)} &= A_k^{l,(i)} - G_k^{l,(i)} (Q_k^{\text{ln}})^T (G_k^{n,(i)} Q_k^n)^{-1} A_k^{n,(i)}, \\
\bar{Q}_k^l &= Q_k^l - (Q_k^{\text{ln}})^T (Q_k^n)^{-1} Q_k^{\text{ln}}.
\end{aligned} \tag{15}$$

x_{k+1}^n is a random variable defined in (10). Here we will estimate its value by using a one-step predictor $x_{k+1|k}^{n,(i)}$ computed in step 2a of the 4.1.4 algorithm for each particle. The terms that differentiate (13)-(15) from the standard Kalman filter one-step prediction equations deal with the decorrelation of the linear and nonlinear state components; see Appendix I in [116] for details.

The conditional linear posterior density $p(x_k^l | x_{0:k}^n, y_{1:k})$ for step k is no longer Gaussian because of its dependence on the likelihood (12) of observing a spike. We, therefore, approximate it using a Gaussian based on the Stochastic State Point Process Filter (SSPPF) methodology developed in [35]. The conditional posterior mean $\hat{x}_{k|k}^{l,(i)}$ and covariance $P_{k|k}^{(i)}$ of the Gaussian approximation become:

$$\begin{aligned}
\hat{x}_{k|k}^{l,(i)} &= \hat{x}_{k|k-1}^{l,(i)} + P_{k|k}^{l,(i)} \left[\frac{\partial \log \lambda_k}{\partial x_k^l} (\Delta N_k - \lambda_k \cdot \Delta) \right]_{\hat{x}_{k|k-1}^{l,(i)}}, \\
(P_{k|k}^{l,(i)})^{-1} &= (P_{k|k-1}^{l,(i)})^{-1} + \left[\left(\frac{\partial \log \lambda_k}{\partial x_k^l} \right)^T (\lambda_k \cdot \Delta) \left(\frac{\partial \log \lambda_k}{\partial x_k^l} \right) - (\Delta N_k - \lambda_k \cdot \Delta) \frac{\partial^2 \log \lambda_k}{\partial x_k^l \partial (x_k^l)^T} \right]_{\hat{x}_{k|k-1}^{l,(i)}}.
\end{aligned} \tag{16}$$

4.1.4 Algorithm

The following is the proposed algorithm.

- 1) Initialization. Initialize each particle's one-step prediction value by sampling the prior distribution: $x_0^{n,(i)} \sim p(x_0^n)$. Set the one-step prediction mean/variance of each particle to the initial values: $\{x_{00}^{l,(i)}, P_{00}^{(i)}\} = \{\bar{x}_0^l, \bar{P}_0\}$.
- 2) Prediction.
 - a) Estimate nonlinear one-step prediction density $p(x_k^n | x_{0:k-1}^n, y_{0:k-1})$ by evaluating the state transition function for each particle x^n in (10).
 - b) Estimate mean/covariance of the linear one-step prediction density $p(x_k^{l,(i)} | x_{0:k}^{n,(i)}, y_{0:k-1})$ for each particle with (13)-(15).
- 3) Evaluate nonlinear importance weights for all $i=1 \dots M$ particles:

$$\tilde{w}_k^{(i)} = w_{k-1}^{(i)} \cdot p(y_k | x_{0:k}^{n,(i)}, y_{0:k-1}).$$
 Normalize the weights: $w_k^{(i)} = \tilde{w}_k^{(i)} / \sum_{j=1}^M \tilde{w}_k^{(j)}$.
- 4) Resample M nonlinear particles (and their corresponding SSPPFs) with replacement so that the probability of particle's selection is proportional to its normalized weight: $\Pr(x_{k|k}^{n,(i)} = x_{k|k-1}^{n,(j)}) \propto w_k^{(j)}$. Set weights $w_k^{(i)} = 1/M$ for $i=1, \dots, M$.
- 5) Use (10) to compute mean/covariance of the conditional linear posterior density for each surviving particle $p(x_k^{l,(i)} | x_{0:k}^{n,(i)}, y_{0:k})$.
- 6) Compute inference functions on linear/nonlinear states. Go to step 2.

Please note that resampling can be skipped for any step; we perform resampling at spike times for all spikes in the recording. Also, any of the available resampling schemes can

be used in step 4 (see [18] for suggestions); we use residual resampling described above. In addition, our use of weight update equation (5) in step 3 means that the proposal density equals the one-step prediction distribution, making our algorithm a bootstrap filter. Also note that the algorithm requires tracking of one linear estimate for each particle, a total of M SSPPFs. Importantly, if steps 2b and 5 are omitted, the MPFPP turns into a regular particle filter with point process observations.

4.2 RESULTS

In this section we apply the estimation procedures to two classes of models. In each case we first use the particle filter with point process observations (PFPP) to estimate the latent state variables from the observed spike trains. We then break up the state variables into nonlinear and linear components and estimate both using the marginalized particle filter with point process observations (MPFPP). We then assess the latent variable estimates in several ways. First, we visually compare estimates to the “true” values used to simulate the spike trains. Second, we quantitatively assess how well our estimation procedure does by computing the total mean squared error (MSE) for both linear and nonlinear estimated traces. Third, we look at the scatter plots of the estimated vs. the actual latent variable values and by compute the coefficient of determination R^2 in order to assess how much of the actual variability in the latent variable traces did our estimates capture. For Model 2 we additionally discuss the implications of applying different algorithms to estimation of the the static maximal conductance model parameters. We compare the MPFPP estimation quality metrics to

those of PFPP and show that the MPFPP accuracy improvements depend on the model type used in estimation.

4.2.1 Model 1

The first model class features a two-dimensional state vector x with one linear and one non-linear component and a direct dependence of the observation equation (18) on both components. The state transition functions have the following form:

$$\begin{aligned} x_{k+1}^n &= 1.5 \arctan(x_k^n) + 0.20x_k^l + \varepsilon_k^n \\ x_{k+1}^l &= 0.98x_k^l + \varepsilon_k^l \end{aligned} \quad (17)$$

where ε s are the following Gaussian noise terms: $\varepsilon_k^n \sim N(0, 0.1)$ and $\varepsilon_k^l \sim N(0, 0.05)$.

The following are the function and matrices from the general form of the algorithm in

$$(8): f_k^n(x_k^n) = 1.5 \arctan(x_k^n), A_k^n(x_k^n) = 0.20, \text{ and } A_k^l(x_k^n) = 0.98.$$

Note that the linear component in (17) is constrained to within the $[-4, 4]$ interval. The *arctan* function forces the nonlinear component into one of the two operating regimes: it either is approximately -2 or is approximately 2. The regime switch occurs when the linear component changes signs, pushing x^n into a different quadrant of the *arctan* function. The state vector components were initialized to values drawn from the following priors: $x_0^n \sim N(0, 1)$ and $x_0^l \sim N(0, 10^{-6})$.

The conditional intensity function linking the observed spike train to the state variables in (17) has the following form:

$$\log(\lambda_k) = \frac{b}{2} (x_k^n)^2 \operatorname{sgn}(x_k^n) + bx_k^l + a, \quad (18)$$

The squaring of the nonlinear term weighs the CIF towards x^n , making it operate in a roughly two-regime mode as well. The coefficients b and a were set to 0.25 and 2.6 respectively to ensure that firing rate ranged between 5 and 30 Hz. The following are the derivatives that stochastic state point process filter (SSPPF) uses in its calculations [35]:

$$\frac{\partial \log \lambda_k}{\partial x_k^l} = b, \quad \frac{\partial^2 \log \lambda_k}{\partial (x_k^l)^2} = 0.$$

The estimation equations for this model are as follows:

$$\begin{aligned} \lambda_k^{(n)} &= \exp\left\{\frac{b}{2}(x_k^n)^2 \operatorname{sgn}(x_k^n) + bx_k^l + \frac{1}{2}b^2 P_{k|k-1}\right\}, \\ \hat{x}_{k+1|k}^{l,(i)} &= 0.98 \hat{x}_{k|k}^{l,(i)} + L_k^{(i)} \left(\hat{x}_{k+1|k}^{n,(i)} - 1.5 \arctan(x_k^n) - 0.2 \hat{x}_{k|k}^{l,(i)} \right), \\ P_{k+1|k}^{(i)} &= 0.98^2 P_{k|k}^{(i)} + 0.05 - \left(L_k^{(i)} \right)^2 N_k^{(i)}, \\ N_k^{(i)} &= 0.2^2 P_{k|k}^{(i)} + 0.1, \\ L_k^{(i)} &= 0.98^2 P_{k|k}^{(i)} / N_k^{(i)}, \\ \hat{x}_{k+1|k+1}^{l,(i)} &= \hat{x}_{k|k-1}^{l,(i)} + b \cdot P_{k+1|k+1}^{l,(i)}, \\ \frac{1}{P_{k+1|k+1}^{l,(i)}} &= \frac{1}{P_{k+1|k}^{l,(i)}} + b^2 (\lambda_k \cdot \Delta). \end{aligned} \tag{19}$$

Note that we chose to use $\hat{x}_{k+1|k}^{n,(i)}$, the one-step prediction estimator for a particle i of the random variable x_{k+1}^n in (10).

Figure 13 displays the nonlinear (top panel) and linear (bottom panel) state estimates along with the "true" value of the state variables used to simulate the spike trains. The left panels display the PFPP estimates; the right show the MPFPP traces. Both algorithms are able to trace the "true" state values reasonably well throughout both positive and negative nonlinear regimes and through some regime transitions. The 95% confidence intervals of the estimates, shown in gray, decrease where the spike rate is

highest for both algorithms, while the lack of spiking increases uncertainty in the estimate. This is observed by comparing the width of the gray area at low spiking (e.g. 100 msec) to the confidence interval size at peak spiking (e.g. 250-300 msec).

The uncertainty of the estimates increased around the nonlinear state variables' regime switches. For instance, the confidence interval of the PFPP nonlinear estimate grew at 370 msec. In this example the uncertainty growth can be attributed to some particles switching from a positive to a negative regime, increasing the overall variance of the marginal nonlinear posterior density and, therefore, the uncertainty of the estimate. Similar situation develops for the MPFPP nonlinear estimate at 100 msec. In both cases, the estimates of the linear state variable was driving the nonlinear regime changes, and therefore, it was expected that a more accurate MPFPP linear estimate has helped reduce the total MPFPP error for whole nonlinear the trajectory.

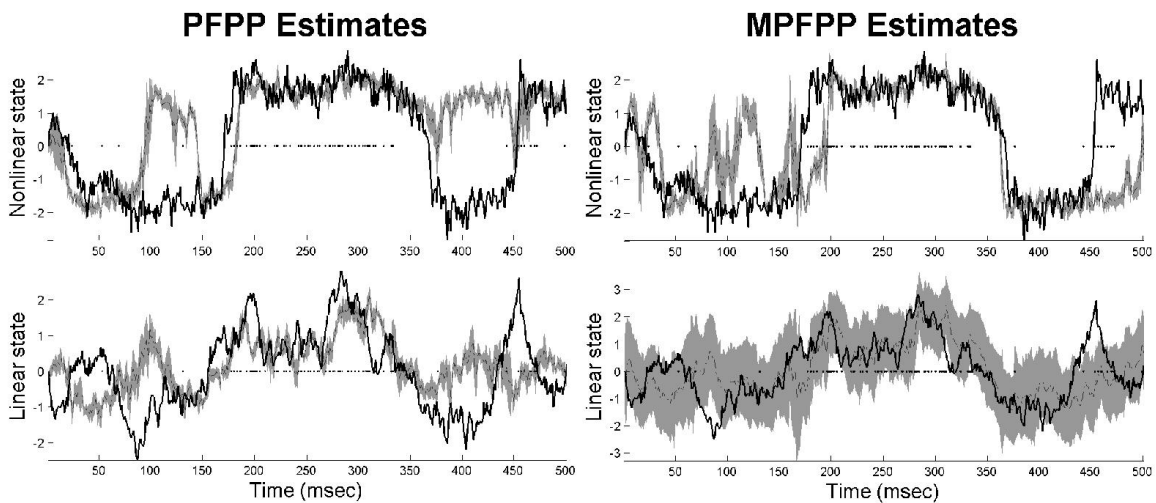


Figure 13. The juxtaposition of the state variables used to simulate the spike trains (actual, black solid line) and their estimates computed from the observed spikes (predicted, gray dash-dot line) using 10 particles. The shaded region surrounding the predicted values represents 95% confidence interval. Spike times are represented by dots at ordinate 0. PFPP stands for particle filter with point process observations; MPFPP is its marginalized variant. Note that the actual nonlinear trace abruptly switches between negative and positive regimes, increasing the difficulty of this dynamic

estimation problem. Neither PFPP nor MPFPP technique is capable of accurately estimating all the switches using such a small number of particles (e.g. see erroneous positive switch by both models at 100 msec). The erroneous regime switches highlight differences in algorithms; for instance, the one at 100 msec is correlated with the upward spikes in linear component estimates. By using random mixture of Gaussians, MPFPP is able to quickly bring the linear estimate in line with the actual trace, while PFPP continues its deflection from the actual trace, sustaining the erroneous positive regime. Similar situation occurs at 350 msec, where PFPP fails to perform the necessary switch because its linear component does not follow the actual trace, but MPFPP switches modes successfully because of better estimated linear component.

The difference between the quality of PFPP and MPFPP linear estimates is better described by plotting the estimated vs. the actual values of the state variables. Bottom panels of Figure 14 show such scatterplots. The bottom right plot shows that MPFPP does a good job at tracking the actual linear state variables throughout the entire range of values. To the contrary, PFPP suffers from not estimating well the values on the range edges; see, for example, the lack of dots in the $[-3, -2]$ corner of the lower left plot. Quantitative assessment of the linear state variables estimation quality corroborates the conclusions above: the correlation between actual and estimated linear variable values for this simulation is 0.75 for MPFPP and 0.61 for PFPP (see Table 5).

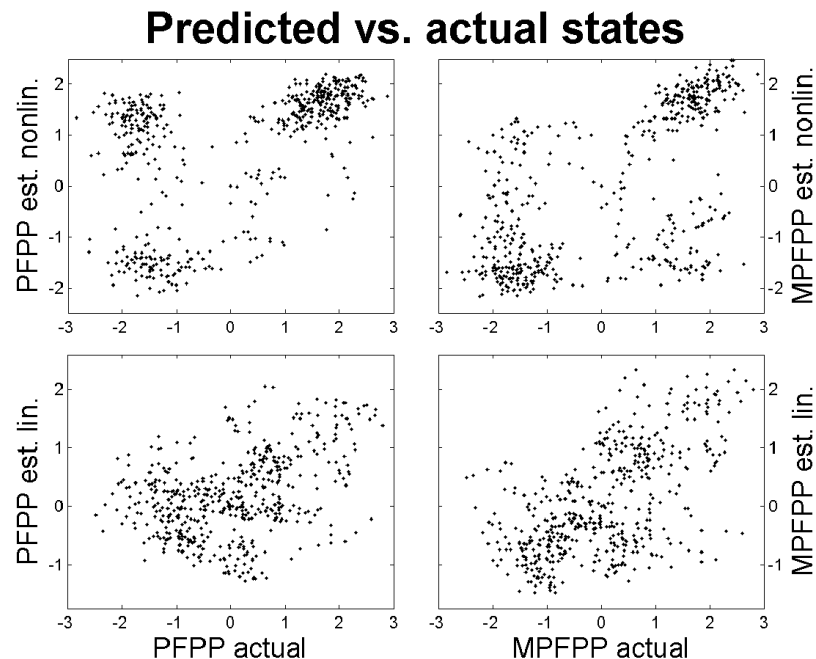


Figure 14. Scatter plots of predicted vs. actual state variable values for Model 1. Two left-hand-side plots feature nonlinear and linear state variable values predicted by the particle filter with point process observations (PFPP); right plots -- by the marginalized particle filter with point process observations (MPFPP). Top plots show nonlinear variables; bottom plots display linear variables. Here, scatter plots visualize the quality of the fit by plotting the values predicted by the model (y-axis, estimated) as a function of the state values used in simulating the observation spike trains (x-axis, actual); a perfect fit lies on a diagonal 45-degree line. The fact that we use a nonlinear model that switches between positive and negative values brings about off-diagonal clustering in the top plots. These clusters indicate a switch in the estimated state trajectory that is not matched by a switch in the simulated state trace. Note that off-diagonal PFPP clusters are more dense than those of MPFPP, suggesting that MPFPP performs smaller number of the mismatched jumps. Also note that the shape of a linear MPFPP point cloud is slimmer and extends more along the diagonal than does the PFPP point cloud. This suggests that MPFPP predictions are better correlated to the actual linear trace than those of PFPP. See text for the discussion of state variable correlation and its quantitative assessment. Nonlinear and linear PFPP correlations are respectively 0.5240 and 0.6072. MPFPP: 0.6381 and 0.7449.

Table 5 shows that marginalized filter does better in nonlinear estimation as well: the MPFPP nonlinear correlation coefficient is higher at 0.64 than PFPP one at 0.52. Top panels of Figure 14 show why: the timely regime switches for the MPFPP estimator are key. The off-diagonal clusters of points are much less pronounced for marginalized filter, while the PFPP cluster centered at $[-2, 1]$ has lead to substantial decrease in the estimation quality.

	ρ		R^2	
	Nonlin.	Linear	Nonlin.	Linear
PFPP	0.524	0.6072	0.2745	0.3687
MPFPP	0.6381	0.7449	0.4072	0.5549

Table 5. Correlation coefficients and coefficient of determination for a single estimation run of Model 1.

So overall, the assessment of the estimated and the actual state variable traces suggests that the marginalized filter excels through a combination of two properties. First, a random mixture of Gaussians is a more accurate estimator of linear state variables than is a small number of particles (10 were used to produce Figures 13-14 and Table 4). Second, a better estimate of the linear component values leads to a more accurate

detection of the nonlinear regime switches, which manifests in a higher accuracy of the nonlinear estimation by the MPFPP.

This superior linear/nonlinear component interplay is highlighted by Figure 15 that displays the joint posterior densities and the estimation particles for a single time step. Left panel displays actual posterior, computed by numerically integrating (4) on a 70-by-70 grid. It features the curved linear/nonlinear relationship, indicative of the *arctan* function in (17). It is also slim in the nonlinear dimension indicating heightened sensitivity of probability mass regions to changes in nonlinear dimension given a specific linear value.

The majority of MPFPP particles are in the negative nonlinear regime, matching the probability mass distribution of the numerically computed posterior. The particles that retained positive nonlinear values align themselves along a "flow corridor" (at the linear value of approximately -1) that will eventually push them into the negative regime. It is important to note that the formation of such "corridors" is possible because a linear value of each particle is in fact a mean of this particle's conditional linear posterior Gaussian density $p(x_k^{l(i)} | x_{0:k}^{n(i)}, y_{1:k})$. These means can be compactly organized along a linear dimension because each mean's variance contributes to the total dispersion of the joint posterior isocontours along the linear dimension. This allows efficient channeling by the MPFPP of the nonlinear values between the positive and negative operating regimes, leading to faster transitions and better overall estimation. Contrast that with the PFPP particles in the left plot, which are much more spread out and are late making the nonlinear regime switch. The fact that PFPP particles are sampled from a two-

dimensional posterior further impedes their efficient travel between regimes, leading to a decrease in performance when compared to the MPFPP, where nonlinear particle space is one-dimensional.

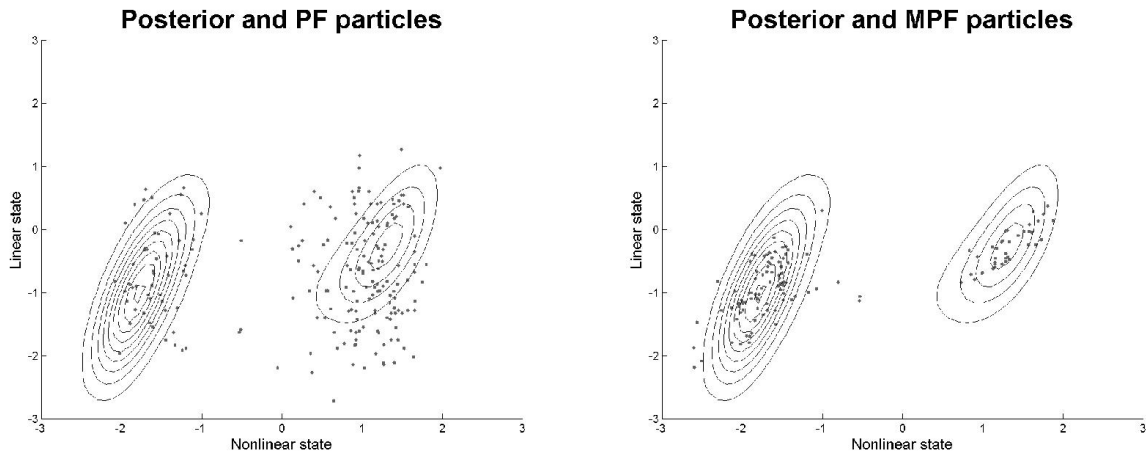


Figure 15. Numerically computed joint posterior densities juxtaposed with the estimation particles for the two state dimensions of Model 1. The plots depict densities at a single time step with no spikes. X-axis represents a nonlinear dimension; y-axis represents a linear dimension. Left panel: PFPP, right panel: MPFPP. Estimates were computed using 100 particles. Notice that the majority of the numerical probability mass has transferred from positive to negative nonlinear regime. MPFPP particles on the left reflect that transition (there are more negative nonlinear particles than positive ones). The majority of PFPP nonlinear particles, however, are still positive, indicating PFPP nonlinear state estimate is lagging behind the transition. Also note the compression of the MPFPP particles when compared to the wider dispersion of PFPP particles in the linear dimension. This occurs because MPFPP particles represent means of Gaussian distribution in the linear dimension; the weighted linear combination of these Gaussians matches the dispersion of the numerically computed joint posterior in the linear dimension.

Figure 16 depicts the comparison of mean squared error across different number of particles used in regular and marginalized particle filters. For each data point we fixed a number of particles and then ran each algorithm on 100 simulated spike trains. We then computed MSE for each state vector dimension across 100 trials and reported its mean along with the 95% confidence interval. We see that for small number of particles, the MPFPP improves the accuracy of the nonlinear component estimation (Figure 16, top plot). This occurs because the same number of particles that cover the entire space

during PFPP estimation need to only cover nonlinear subspace during MPFPP estimation. The linear component estimation error (Figure 15, bottom panel) also decreases during MPFPP estimation. The error decrease is larger for small number of particles, indicating that the random mixture of a small number of Gaussians is a substantially better estimator of state linear variables than a small set of random samples. The nonlinear improvement is relatively constant across the number of particles, indicating that the error reduction comes from the same mechanism, likely the reduction of dimensionality by the MPFPP.

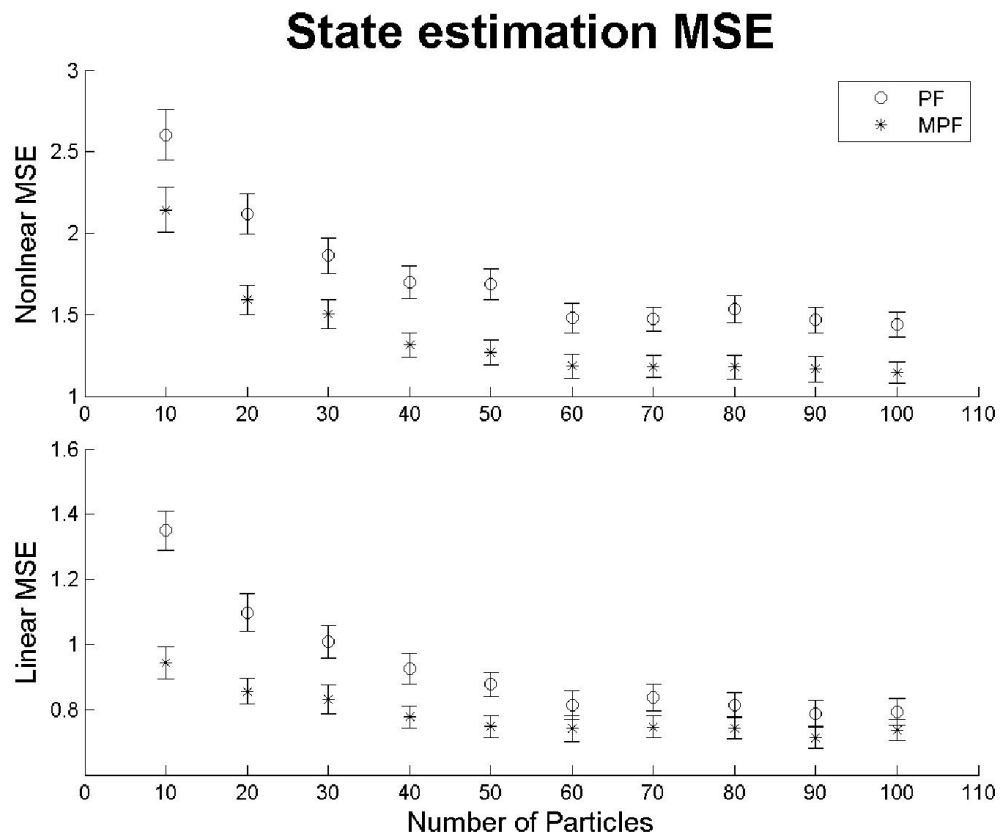


Figure 16. The average mean square error (MSE) taken across 100 estimation runs of Model 1 plotted as a function of the number of particles used in estimation. The top plot lists nonlinear component errors; bottom plot shows linear errors. Error bars represent 95% confidence intervals. Note that the estimation accuracy for both regular (PFPP) and marginalized (MPFPP) particle filters increases as the number of particles grows; this is expected, as a large sample size provides a better estimate of the posterior distribution than the small sample size. Also note that the difference in accuracy between nonlinear PFPP and nonlinear MPFPP prediction remains roughly constant as

number of particles rises; however it shrinks for linear predictions. Indeed, the bottom plot shows that the MPFPP linear prediction for 10 particles is more accurate than the PFPP linear prediction for 20 particles. Compare that to the top plot, where a 10-particle nonlinear MPF prediction is as accurate as a 20-particle nonlinear PF prediction. This suggests MPF as a tool of choice for estimation problems where the allowable computational load (through the total number of particles or total system dimensionality) is restricted, and where linear components are of primary interest. Moreover, if one is to choose MPFPP with the smaller number of particles over PFPP with the larger number of particles, the loss of nonlinear accuracy due to particle reduction is mitigated by the smaller MPF estimation error.

The plot also highlights the utility of the MPFPP estimation when compared to regular particle filters. MPFPP with only 20 particles does better than PFPP with 30 and 40 particles. Only when the number of PFPP particles reaches 50, do we observe similar performance. This suggests that MPFPP can be readily substituted for PFPP, leading to shorter estimator running times without loss of accuracy.

4.2.2 Model 2

We also used a realistic biophysical system to test the marginalized particle filter with point process observations (MPFPP) algorithm. We chose Hodgkin-Huxley system because of its ability to accurately describe the dynamics of the membrane potential for a large variety of neuron types. The original system is expressed as a set of coupled differential equations [57]:

$$\begin{aligned} C\dot{V} &= I - g_K n^4 (V - E_K) - g_{Na} m^3 h (V - E_{Na}) - g_L (V - E_L) \\ \dot{n} &= \alpha_n(V)(1-n) - \beta_n(V)n \\ \dot{m} &= \alpha_m(V_k)(1-m_k) - \beta_m(V_k)m_k \\ \dot{h} &= \alpha_h(V)(1-h) - \beta_h(V)h \end{aligned}$$

where

$$\begin{aligned}
\alpha_n(V) &= 0.01 \frac{10-V}{\exp\left(\frac{10-V}{10}\right)-1}, & \beta_n(V) &= 0.125 \exp\left(\frac{-V}{80}\right), \\
\alpha_m(V) &= 0.1 \frac{25-V}{\exp\left(\frac{25-V}{10}\right)-1}, & \beta_m(V) &= 4 \exp\left(\frac{-V}{18}\right), \\
\alpha_h(V) &= 0.07 \exp\left(\frac{-V}{20}\right), & \beta_h(V) &= \frac{1}{\exp\left(\frac{30-V}{10}\right)+1}.
\end{aligned} \tag{21}$$

The equations describe the propagation of current through the neuron as a function of membrane potential V , membrane capacitance C , and three different ion currents: sodium (Na), potassium (K), and leak (L) currents. I represents applied current. Each ion channel may be characterized by its maximal conductance g . The sodium and potassium channels are modeled explicitly, and their time- and voltage-dependent behavior is specified by the gating variables n , m , and h . Leak channel is simplified by only having a voltage-independent maximal conductance variable g_L .

The following parameter values from [57] were used both to simulate and to estimate the latent variables: $E_K=-12$ mV, $E_{Na}=120$ mV, $E_L=10.6$ mV, $I=10$ $\mu\text{A}/\text{cm}^2$, $C=1.0$ $\mu\text{F}/\text{cm}^2$. The maximal conductance of the leak channel was kept constant at $g_L=0.3$ mS/cm². The maximal conductances of the voltage-dependent channels were simulated as a random walk (22) with initial values $g_{K,0}=36$ mS/cm² and $g_{Na,0}=120$ mS/cm² [57]. Please note that if any parameter values are not explicitly mentioned here, they are exactly as in [57].

In order to test the MPFPP algorithm, we first simulated the trajectory of the membrane potential for 300 msec. In order to avoid model misspecification, we

simulated and estimated the state trajectories using a standard particle filter technique of combining the static parameters and latent variables outlined in [80]. Specifically, we varied the static parameters a little with time and appended them to the state variables, producing a six-dimensional state latent vector $[V, n, m, h, g_K, g_{Na}]^T$. The following discretized version of the Hodgkin-Huxley system was used to simulate the membrane potential variables as well maximal conductance model parameters:

$$\begin{bmatrix} V_{k+1} \\ n_{k+1} \\ m_{k+1} \\ h_{k+1} \\ g_{K,k+1} \\ g_{Na,k+1} \end{bmatrix} = \begin{bmatrix} V_k + (I - g_K n_k^4 (V_k - E_K) - g_{Na,k} m_k^3 h_k (V_k - E_{Na}) - g_L (V_k - E_L)) \cdot \Delta / C + \varepsilon_k^V \\ n_k + (\alpha_n(V_k)(1 - n_k) - \beta_n(V_k)n_k)\Delta + \varepsilon_k^h \\ m_k + (\alpha_m(V_k)(1 - m_k) - \beta_m(V_k)m_k)\Delta + \varepsilon_k^h \\ h_k + (\alpha_h(V_k)(1 - h_k) - \beta_h(V_k)h_k)\Delta + \varepsilon_k^h \\ g_{K,k} + \varepsilon_k^K \\ g_{Na,k} + \varepsilon_k^{Na} \end{bmatrix} \quad (22)$$

The step size was set to $\Delta=0.2$ msec. ε 's are uncorrelated noise terms with the following variances: $\sigma_V^2 = 0.005^2 \cdot \Delta$, $\sigma_n^2 = \sigma_m^2 = \sigma_h^2 = 10^{-100}$, and $\sigma_K^2 = \sigma_{Na}^2 = 0.005^2 \cdot \Delta$. α s and β s remained as in (21) above. Noise terms with extremely small variances were added to the gating variables n , m , and h in order to ensure that covariance matrices remain full rank throughout the estimation procedure. Noise variances of other variables were selected to ensure that variability did not interfere with the dynamical system and that consistent spiking was maintained for the majority of simulated trajectory.

After the membrane potential trace was simulated, we identified the trajectory's depolarization peaks and designated those as spike times, producing spike train observations.

Upon simulating the state and observation trajectories, we applied the PFPP algorithm to estimate the latent state variables' trace from the generated spike train observations. The system in (22) was used to model the temporal evolution of state variables with the following exception: the variance of the maximal conductance noise terms were increased to $\sigma_K^2 = \sigma_{Na}^2 = 0.5^2 \cdot \Delta$.

In order to relate the continuous state to the spike train observation variables, we defined a conditional intensity function as

$$\lambda_k(V_{0:k}, V_{k+1:k+m}) = \sum_{\tau=1}^{k+m} [g(V_\tau) \cdot f(\tau - k)], \quad (23)$$

$$g(x) = \eta \cdot \frac{\exp(v(x - V_{th}))}{1 + \exp(v(x - V_{th}))}, \quad f(x) = p^{|x|}.$$

The following parameter values are used (from [89]): $\eta=1.622$, $v=0.1$, $p=0.9$, $V_{th}=80$ mV.

This formulation of the spike probability model depends on the past as well as future membrane potential values. Because we only resample particles at spike times, this probability model increases the spike likelihood for particles that have recently spiked or will spike in the near future, smoothing the likelihood landscape and increasing relevant particles' survival probability during resampling. For more details on this spike likelihood model formulation see section 2.3 of [89].

Here, the sigmoid function $g(\cdot)$ reflects the effect of increasing voltage values on the spike intensity, in which parameter v controls the growth rate of the function and parameter η sets the upper limit; V_{th} is a voltage threshold, above which the probability of spiking rapidly increases. The function $f(\cdot)$ is a weighting function, which measures the effect of the voltage at time bins away from time bin k ; p is a constant less than 1.

Each particle was initialized as follows: initial membrane potential V_0 was sampled from a uniform distribution on an interval $[-5, 55]$ mV. The activation parameter initial value n_0 was computed from V_0 using the following “resting state” equation ([57], page 507): $n_0 = \alpha_{n_0} / (\alpha_{n_0} + \beta_{n_0})$, where α_n and β_n are specified by (21). Initial values m_0 and h_0 were calculated similarly (see Eq. 17, 18 in [57]). The initial conductance values g_K and g_{Na} were uniformly distributed on an interval $[0, 100]$ and $[0, 300]$ mS/cm² respectively. The estimation then proceeded as described in section 4.1.1 and 4.1.4.

We then estimated the latent variable trajectories using MPFPP and compared the results to those obtained with the regular particle filter. To apply MPFPP, we separated the difference equations into a linear and a nonlinear dimensions and put them in the general form specified by (10), where $x^n = [V, n, m, h]^T$ is a nonlinear and $x^l = [g_K, g_{Na}]^T$ is a linear component. We restate the system in (22) as:

$$\begin{aligned} \begin{bmatrix} V_{k+1} \\ n_{k+1} \\ m_{k+1} \\ h_{k+1} \end{bmatrix} &= \begin{bmatrix} V_k + (I - g_L(V_k - E_L)) \cdot \Delta / C \\ n_k + (\alpha_n(V_k)(1 - n_k) - \beta_n(V_k)n_k)\Delta \\ m_k + (\alpha_m(V_k)(1 - m_k) - \beta_m(V_k)m_k)\Delta \\ h_k + (\alpha_h(V_k)(1 - h_k) - \beta_h(V_k)h_k)\Delta \end{bmatrix} + A_k^n(x_k^n) \begin{bmatrix} g_{K,k} \\ g_{Na,k} \end{bmatrix} + \varepsilon_k^n, \\ \begin{bmatrix} g_{K,k+1} \\ g_{Na,k+1} \end{bmatrix} &= \begin{bmatrix} 0 \\ 0 \end{bmatrix} + A_k^l \begin{bmatrix} g_{K,k} \\ g_{Na,k} \end{bmatrix} + \varepsilon_k^l. \end{aligned} \quad (24)$$

Here, the appropriate functions and matrices from (10) are as follows:

$$\begin{aligned}
f_k^n(x_k^n) &= \begin{bmatrix} V_k + (I - g_L(V_k - E_L)) \cdot \Delta / C \\ n_k + (\alpha_n(V_k)(1 - n_k) - \beta_n(V_k)n_k)\Delta \\ m_k + (\alpha_m(V_k)(1 - m_k) - \beta_m(V_k)m_k)\Delta \\ h_k + (\alpha_h(V_k)(1 - h_k) - \beta_h(V_k)h_k)\Delta \end{bmatrix}, \\
f_k^l(x_k^n) &= \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \\
A_k^n(x_k^n) &= \begin{bmatrix} -g_{K,k}n_k^4(V_k - E_K) & -g_{Na,k}m_k^3h_k(V_k - E_{Na}) \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \frac{\Delta}{C}, \\
\text{and } A_k^l &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.
\end{aligned} \tag{25}$$

The uncorrelated nonlinear and linear noise terms are

$$\begin{aligned}
\varepsilon_k^n &\sim N(\underline{0}, Q^n) = N(\underline{0}, \text{diag}(\sigma_V^2, \sigma_n^2, \sigma_m^2, \sigma_h^2)), \\
\varepsilon_k^l &\sim N(\underline{0}, Q^l) = N(\underline{0}, \text{diag}(\sigma_K^2, \sigma_{Na}^2)).
\end{aligned} \tag{26}$$

The parameter values, spike probability model, and initialization procedures were identical to the PFPP ones described above. The prior distribution on the linear component was a Gaussian with the following mean/covariance:

$$\{x_{00}^{l,(i)}, P_{00}^{(i)}\} = \left\{ \begin{bmatrix} g_{K,0} \\ g_{Na,0} \end{bmatrix}, \begin{bmatrix} 2000 & 0 \\ 0 & 2000 \end{bmatrix} \right\}.$$

We used the spike probability model (23) to relate spike train observations to MPFPP state components. Note that because of its dependence solely on membrane potential V , this formulation of the spike probability model fits into the general

framework (12) with $h_k(x_k^n) = \log \sum_{\tau=1}^{k+m} \left[\eta \cdot \frac{\exp(v(V_\tau - V_{th}))}{1 + \exp(v(V_\tau - V_{th}))} \cdot p^{|\tau-k|} \right]$ and $C_k = 0$.

Algorithm defined in section 4.1.4 has then been applied. Because the spike probability model (23) depends on voltage only, the derivatives of its logarithm with respect to the linear component are 0. This allows us to write both the time update and observation correction steps as a single set of equations:

$$\begin{aligned}
\mathbf{x}_{k+1|k+1}^{l,(i)} &= \mathbf{x}_{k|k}^{l,(i)} + P_{k|k}^{(i)} \left(A^{n,(i)} \right)^T \cdot \left(N^{(i)} \right)^{-1} \cdot \left(z_k^{(i)} - A^{n,(i)} \mathbf{x}_{k|k}^{l,(i)} \right), \\
P_{k+1|k+1}^{(i)} &= P_{k|k}^{(i)} + Q^l - L^{(i)} N^{(i)} \left(L^{(i)} \right)^T, \\
N^{(i)} &= A^{n,(i)} P_{k|k}^{(i)} \left(A^{n,(i)} \right)^T + Q^n, \\
L^{(i)} &= P_{k|k}^{(i)} \left(A^{n,(i)} \right)^T \left(N^{(i)} \right)^{-1}, \\
z_k^{(i)} &= \mathbf{x}_{k+1}^n - f^n \left(\mathbf{x}_{k|k}^{n,(i)} \right),
\end{aligned} \tag{27}$$

where Q^n is defined in (26). Note that we chose to use $\hat{\mathbf{x}}_{k+1|k}^{n,(i)}$, the one-step prediction estimator for a particle i of the value \mathbf{x}_{k+1}^n in (19). Also, the computation of weights requires computing a new rate $\lambda_{k+1}^{(n)}$. Because $C_k = 0$, we use the rate $\lambda_{k+1}^{(n)} = \lambda_{k+1}$ from (23) in MPFPP weight computations.

The PFPP and MPFPP compute estimates of both linear and nonlinear state variables. Here we are primarily interested in the membrane potential and maximal conductance estimates because the former combines the effects that all state variables have on spike production, and the latter describe the biophysical properties of the ion channels. Figure 17 juxtaposes the simulated and the estimated traces of these variables. Both PFPP and MPFPP estimates reproduce the general outline of the voltage trace, including periods of depolarization that constitute action potentials (Figure 17, top plot). Note that initially spikes are estimated prematurely, indicating a large number of particles spiking early. Subsequently, as maximal conductance values move closer to the

simulated ones, the large premature spikes disappear. However, the variability in depolarization timing is not gone entirely; Figure 18 (top plot) zoom in on the last 50 msec of the recording and depicts the remaining estimation discrepancies. Despite these discrepancies, the 500-particle estimator captured an average of 47% (S.E.=1.8%) of the variance of the simulated voltage across 100 runs. Note also that, by eliminating particles that are unlikely to spike in the near future, resampling at spike times eliminates false positive action potentials that are late relative to the actual spikes and increases the overall estimation accuracy.

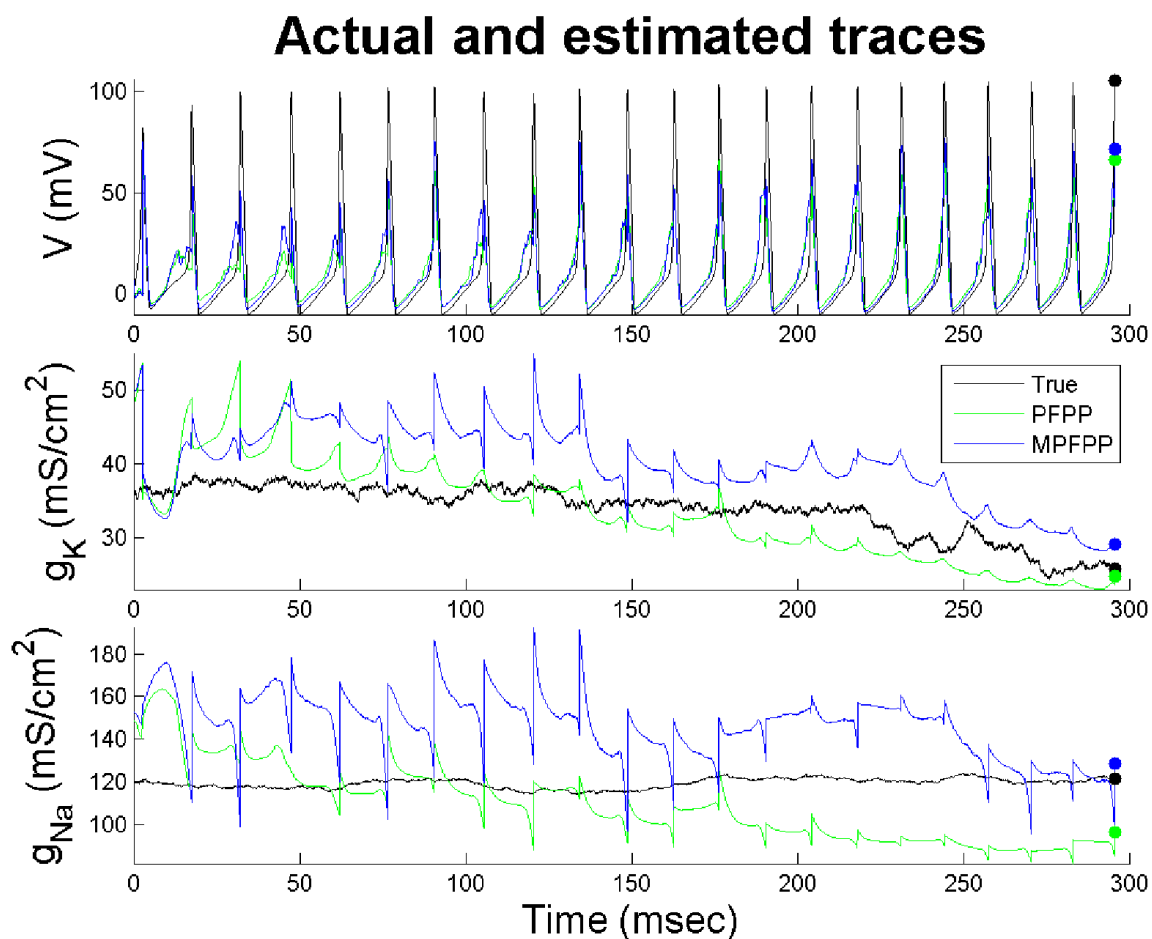


Figure 17. The estimates of membrane potential (top plot) and maximal conductance values (two bottom plots) computed by regular (green) and marginalized (blue) particle filter algorithms using

500 particles. Black trace denotes “true” simulated voltage. Large dots indicate the final estimate at the time of the last spike. Note that both PFPP and MPFPP track the general contour of the simulated voltage trace. Both estimates fail to accurately predict sharp depolarizations forming the action potentials, making the final predictions biased downwards (see voltage dots). Both algorithms trace the overall maximal conductance trends (downwards for g_K and horizontal for g_{Na}), but, as expected, fail to accurately decode the “true” values for each individual time step. Also note that spikes seem to carry more information for MPFPP than for PFPP as indicated by larger magnitude jumps in the blue trace at spike times.

The bottom plot of Figure 17 shows the estimated trajectories of maximal conductances g_K and g_{Na} (green and blue). They follow the overall trend of the simulated conductance trace (black), but fail to faithfully reproduce the values at individual time points. This occurs for two reasons. First, the confidence intervals of the MPFPP conductance estimators are wide. For instance, the variance of the g_{Na} parameter estimated by MPFPP at the time of the last spike (blue dot in Figures 17, 18) is 16.01, while the sample variance of the entire simulated g_{Na} trace is 5.33. Given the large uncertainty about the estimate, we do not expect it to capture the local variability of the simulated trace well. Second, the value of the estimate changes drastically during resampling, increasing the overall variance of the estimate. For example, the estimated trace shown in Figure 18 captures an average of $R_K^2 = 23\%$ (S.E.=2.2%) and $R_{Na}^2 = 18\%$ (S.E.=1.9%) of total variance of the simulated trace for the same period (averages are taken across the 100 estimation runs). None of this is surprising because we only use 500 particles in this simulation experiment (compare that to 10,000 used in [89]).

It is also worth noting that our primary interest here lies not in capturing the step-to-step conductance variability, but in convergence of the overall estimate to the true values. Figure 17 shows that this is the case, with the overall trends of the simulated conductance trends being followed. Better trend estimation with tighter confidence

intervals are possible with increasing the number of particles used to simulate the state variables and with increasing the length of the simulations.

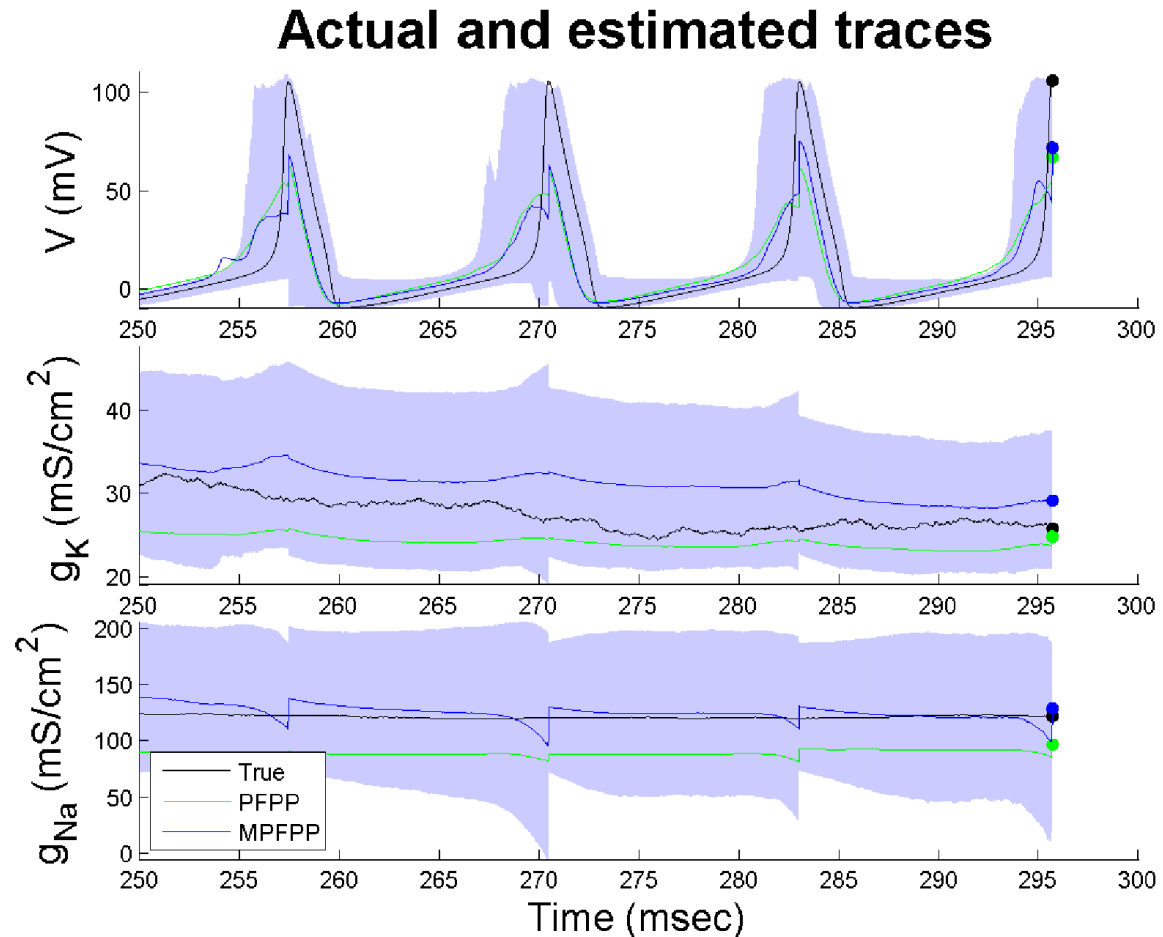


Figure 18. A detailed view of the trailing 50 milliseconds of the estimation traces from Figure 17. The shaded blue areas are MPFPP 95% confidence intervals; the remaining traces are as in Figure 17. Note that while both PFPP and MPFPP capture the voltage increases during action potentials, the estimates (blue, green traces) fail to account for the very tops of each peak. However, the peaks of the action potentials are within the MPFPP estimator's 95% confidence intervals. Also note that some action potentials are estimated to occur early, but none are estimated late. In addition, the local variability of maximal conductances g_K and g_{Na} is not captured by the estimates. For example, simulated g_K trace is convex between 260 and 270 msec, while the estimates are concave.

Figure 19 depicts the estimates of the maximal conductance parameters in the Hodgkin-Huxley model along with the particles that were used to compute those estimates. Note that the “true” $(g_{Na,0}, g_{K,0})$ conductance value (22) lies within the

estimator confidence bounds. Also note that the 95% confidence interval for the MPFPP estimator is slimmer and lies along the constant-spike-rate manifold (black trace in Figure 19).

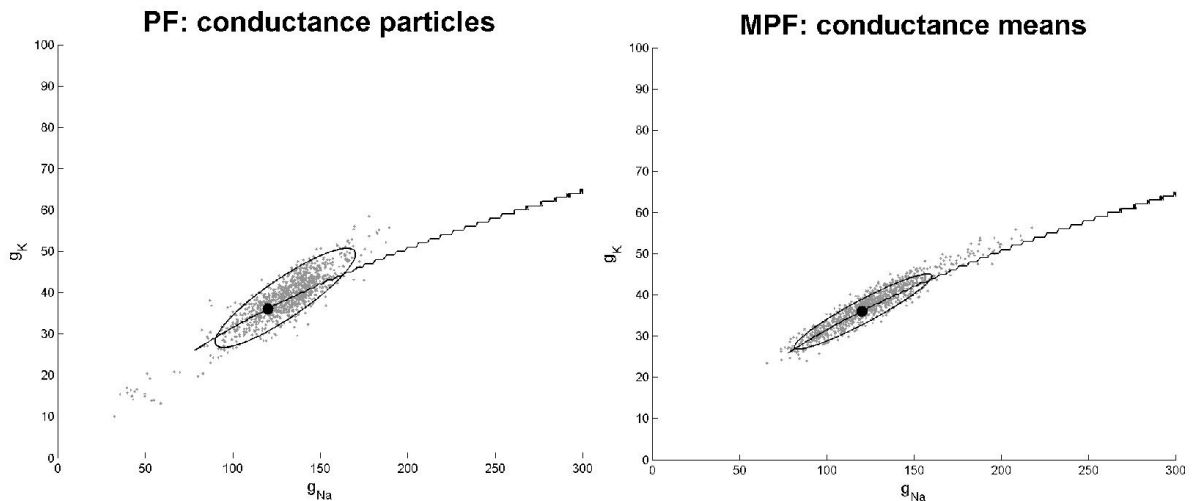


Figure 19. Estimates of maximum conductance parameters for the Hodgkin-Huxley model. Each dot in the left panel indicates a position of a single PF particle in the g_{Na}/g_K space after a 300 msec recording is processed. Each dot in the right panel represents an estimate of where a linear components of the state vector (g_K and g_{Na} in this case) lie in the g_{Na}/g_K space. Large black dot is a true “experimental” value used to simulate the data. The black diagonal curves represents an estimate of a manifold within the g_{Na}/g_K space that produces the same spiking rate that we get with g_{Na} and g_K values used to simulate the model.

We also compare the estimation accuracy of the regular and marginalized particle filters using the root mean squared error (RMSE) measure between the state variable values used to simulate the spike trains and the values that were estimated from the observed spiking. For the cross-membrane potential and maximal conductance trajectories, this corresponds to comparing traces in Figure 17 (top plot).

Figure 20 shows the RMSE values and highlights several important points. First, both algorithms do remarkably well given the complexity of the underlying state dynamical system and the limited amount of information that binary spike train provides

about the continuous state variables. Second, MPFPP voltage estimation does significantly better than its PFPP estimation for both 100 and 300 particles. In fact, the 100-particle MPFPP estimate does not significantly differ from the 300-particle PFPP estimate. This was surprising given the complexity and the structure of the system. And third, the RMSE values do not inform us very much about the difference in quality of the maximal conductance estimation between algorithms. Specifically, for g_K estimates there is no statistical difference between PFPP and MPFPP errors for any number of particles we used. MPFPP estimator of g_{Na} does better than PFPP estimator for 100 particles, but fails to improve at 300 and 500 particles.

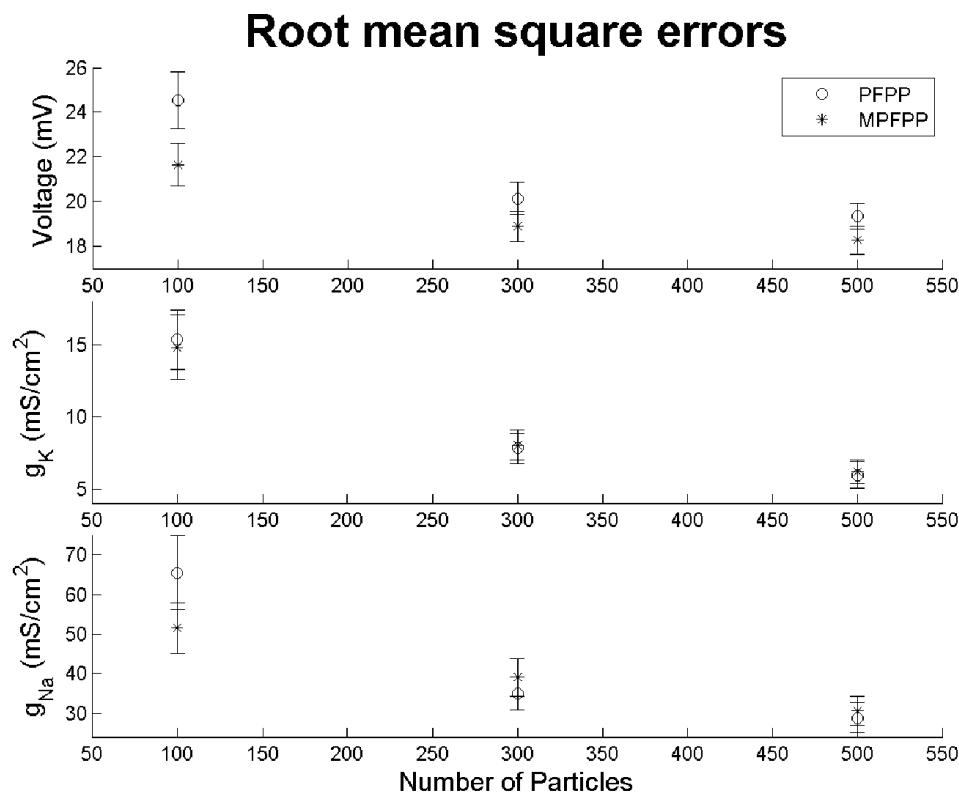


Figure 20. Root mean square error (RMSE) measures for cross-membrane potential and maximal conductance estimates computed by PFPP and MPFPP estimators for Model 2. RMSEs are computed across 100 runs and shown as a function of a number of particles used during each estimation. Top plot depicts RMSE values between the simulated and the estimated cross-membrane

potential traces. Middle plot shows RMSEs between simulated and estimated potassium maximum conductance values. Bottom plot shows the same for sodium maximal conductance. Note that MPFPP does better for the cross-membrane potential for any number of particles (top plot). For 100 and 300 particles, the voltage RMSE difference is statistically significant. To the contrary, the conductance estimation quality is similar between the two estimators. For example, the 100-particle MPFPP estimate of g_{Na} is better than that by PFPP estimate, but 100-particle g_K estimates are statistically indistinguishable (and continue to be for any number of particles).

While the estimation quality of the membrane potential dynamics is important, the estimation of the static values of the maximal conductance parameters of the model constitutes a more physiologically relevant problem. Here, we used g_K and g_{Na} values computed by (7) at the last spike time of the dataset as a static parameter estimate. The time of the last spike was chosen because resampling improves the estimation quality at that point (e.g. see the adjustment of voltage trace at every spike time in Figure 17). Figure 21 shows the distributions of the static estimates as a function of the number of particles used in estimation.

The juxtaposition of PFPP (green) and MPFPP (blue) static estimate distributions in Figure 21 reveals that their structure is different for the two estimators. The MPFPP distributions are often more dispersed (e.g. g_{Na} 300- and 500-particle distributions). However, they are also more consistent. In particular, the parameter values estimated by Hodgkin and Huxley in their original paper ($g_{K,0}=36$ mS/cm² and $g_{Na,0}=120$ mS/cm²) [57] and that were used to initialize the simulation (black horizontal line in Figure 21) has always ended up within the 95% confidence interval of the MPFPP estimate's median (the notch in the box plot). That is not the case for the distributions of PFPP estimates. This indicates that the MPFPP estimates are more reliable on average than the PFPP ones are, but that they can also sometimes produce values that are wildly different from the "true" simulated values.

Also note that the final voltage estimates are biased downwards for both algorithms. This suggests the inability of the two estimators to correctly represent the fast depolarization of action potentials. Figures 17 and 18 show these undershoots in more detail. It is worth mentioning, however, that the action potential peaks are contained within the 95% confidence intervals of both estimators (see Figure 18).

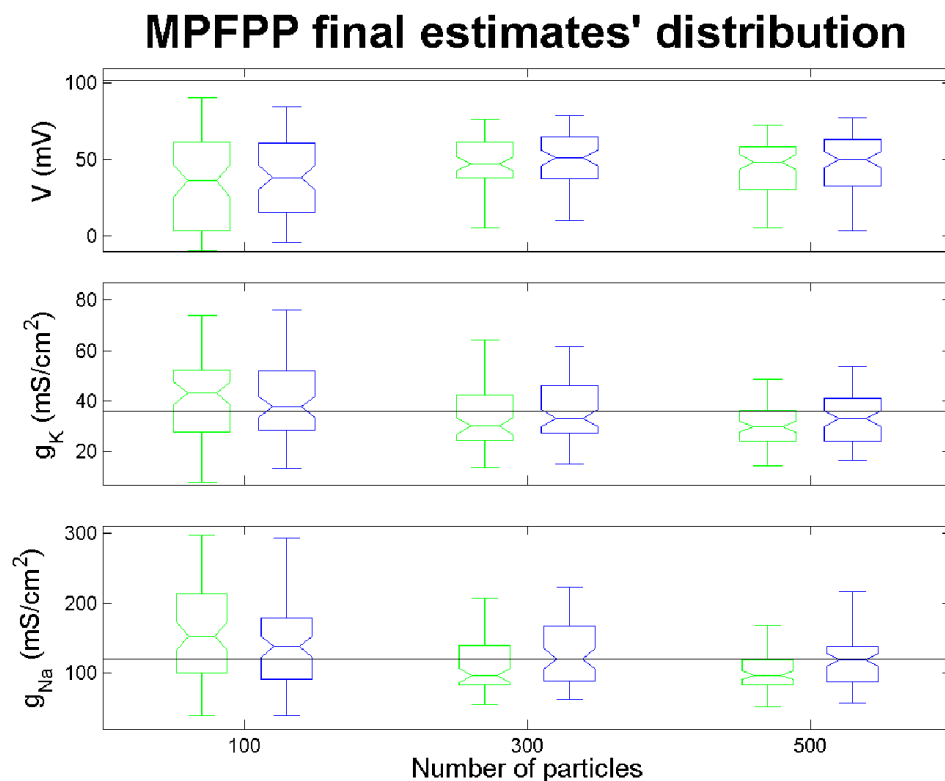


Figure 21. Box plots of the voltage and maximal conductance estimates taken at the time of the last spike in the recording (green denotes PFPP; blue marks MPFPP). Distributions are taken over 100 estimation runs. Box plot whiskers denote 2.5 and 97.5 percentiles; top and bottom box boundaries represent 75 and 25 percentiles respectively; horizontal line in the middle of the box is the sample median, and notches represent its 95% confidence interval. Black horizontal lines are mean simulated values for each variable. MPFPP sample median for maximal conductances is a better estimator of "true" values (black lines): it's perfect for 300 and 500 g_{Na} particles and within sample median's 95% confidence for all others. PFPP estimates of both g_K and g_{Na} is shifted for all particles. Note that voltage distributions are shifted downwards, indicating underestimation of the sharp peak of action potentials by both algorithms.

Last, but not least, it is important to note that a model of spike production during simulation and a spike likelihood model used during the state estimation do not match for Model 2. Specifically, no explicit probability model was used for simulating the spike trains: spike times were generated by noting where the peaks of the action potentials generated by the voltage equation (22) lie. During estimation, however, an explicit spike probability model (23) was used. We do not expect this model misspecification to substantially affect the results we reported because (23) affords high probability of spike occurrence only to voltage values higher than $V_{th}=80$ mV, which is at the very top of the depolarization voltage peak. Using (23) does increase variability of spike time relative to the very top of the action potential peak, but because we weigh this probability by the past and future voltage values, the effect of this variability on the particle weights should remain minimal.

4.2.3 Conclusions

Overall, we were able to show that MPFPP estimator is on average more accurate than PFPP for two types of models and a variety of goodness-of-fit assessment metrics. Specifically, the observations in Model 1 depend on both linear and nonlinear components, and for this model class the increase in MPFPP linear and nonlinear accuracy is obtained for all particle amounts that we tested. In contrast, the observation model (23) of the Hodgkin-Huxley system is independent of the linear components, and we were not able to obtain the decrease in linear RMSE error we saw in Model 1. But, using the distribution of linear static estimates across multiple estimation runs, we showed that MPFPP for that model type is more consistent in estimating the mean

parameter values used in data simulation. This suggests that the MPFPP algorithm is robust to model classes and constitutes a powerful tool for employing spike trains recorded from individual neurons to study these neurons' biophysical properties.

4.3 DISCUSSION

The primary purpose of this chapter is to expand the point process state-space filtering framework by developing and testing a joint linear/nonlinear dynamic estimation algorithm called marginalized particle filter with point process observations (MPFPP). Above we developed the MPFPP theory and compared the algorithm performance to the standard particle filter with point process observations (PFPP) on two classes of models. First, we used a model with a direct dependence of observations on the linear state variables and showed that MPFPP estimates both linear and nonlinear variables more accurately than PFPP does in this setting. Then we applied MPFPP to a biophysical Hodgkin-Huxley model where observations were independent of the linear component. We demonstrated that MPFPP is well-suited for the simultaneous estimation of state variables and model parameters and that it is more accurate than PFPP in estimating the nonlinear components for this model class. We also used simulations to show that an MPFPP linear variable estimates taken at a last spike time are more consistent than those produced by PFPP.

Marginalization is not a new idea; it has previously been applied to dynamic estimation of continuous signals. In particular, Schön et al. showed that marginalized particle filter for continuous observations increased the estimation accuracy of state variables [50]. The authors used observation model that is independent of linear state

components to decrease the linear root mean square error (RMSE) (nonlinear one remained the same). Using marginalization, Gustafsson et al. [49] improved the estimation accuracy of both linear and nonlinear state variables; he used a nonlinear-only observation model. Our Model 2, which also features nonlinear-only observation model dependence, have failed to improve the linear estimate accuracy (Figure 19), but did reduce RMSE for the nonlinear component. The reason is likely a dependence of our observation model on a lone nonlinear variable out of four available. This setup accurately represents the likelihood of spike production because in real neurons the membrane potential is the primary observed signal that aggregates the effects of the multiple ion channels represented by the other three gating variables in the Hodgkin-Huxley dynamical system. But such setup might impede the MPFPP from utilizing its advantages to the fullest as we discuss in detail below.

The estimation accuracy improvement that Model 2 showed differed from that of Model 1 in the following manner: Model 2 only bettered the nonlinear estimates (Figure 19), while Model 1 showed a statistically significant decrease in the RMSE measures for both linear and nonlinear state variables for different amounts of particles (Figure 16). The reason for this is likely two-fold. First, Model 2's $A_k^n(x_k^n)$ in (25) has row rank 1, eliminating the direct influence of the nonlinear gating variables n, m, h on the linear components and leaving voltage as the only nonlinear variables directly contributing to the estimates $\hat{x}_k^{l,(i)}$ of the linear mean. The voltage-only influence limits the mobility of the linear mean estimators in the g_K/g_{Na} plane to trajectories perpendicular to the constant-spike-rate manifold in Figure 18. In the absence of resampling, this prevents

mean estimators $\hat{x}_k^{l,(i)}$ from effectively covering all of the manifold, increases particle clustering, and decreases the linear estimation accuracy. Model 1, in contrast, does not have this problem as $A_k^n(x_k^n)$ is a constant.

The second discrepancy between the Models 1 and 2 that likely contributes to the accuracy of the linear component estimation is the form of the observation probability model (11). Conditional intensity function (CIF) for Model 1 (18) retains both x^l and x^n terms. This ensures that linear posterior estimates in (16) are adjusted up the spike rate gradient at every time step. Instantaneous spiking rate for the Hodgkin-Huxley model (23) lacks linear term dependence and only includes on membrane potential V (23). This eliminates the posterior adjustment of the linear estimates, decreases the amount nonlinear information that linear estimator utilizes, and diminishes its accuracy.

Estimation accuracy improvements shown in Figure 16 make a compelling case for applying MPFPP in situations where a limitation exists on the computational resources. For instance, applying MPFPP can reduce the number of particles required to accurately estimate the state variables, reducing computation time. In most situations, a reduction in the number of computations for the same number of particles is also possible [75]. This is very useful for complex biophysical systems because substantial number of particles is usually required to dynamically track the unobserved state variables even for simple all-linear Gaussian random walk dynamics: 1000 particles for a 7-dimensional system [134], 2500 particles for a 6-dimensional system [13], and 10000 particles for a 10-dimensional system [37]. However, the formulation of the state transition and observation models become extremely important. For instance, the observation equation

must depend on both linear and nonlinear terms to increase estimation accuracy for all state variables. In general, our results suggest that MPFPP exhibits a trade-off between model complexity and estimation accuracy. It manifests itself here as an increase in the linear RMSE resulting from a decrease in the observation model complexity (by dropping a linear term). However, because marginalized particle filter for continuous observations does not seem to exhibit such trade-off [49, 50], further work is necessary to verify that this is something indicative of MPFPP that was brought on by using point observations.

The power and utility of MPFPP and care that needs to be taken in its application show the direction for future work. For example, we can use a different linear/nonlinear subdivision for the Hodgkin-Huxley model. Instead of limiting the linear components to only sodium and potassium maximal conductances we can include maximal leak conductance as well as some or all of the Nernst equilibrium potentials into the linear components. This would change the linear/nonlinear interdependence within the model that we believe has contributed to our inability to improve the nonlinear dimension accuracy in Model 2, and increase our chances of success. Alternatively, we can use a different type of neuronal current model (e.g. FitzHugh-Nagumo and its extensions) to see if we can improve accuracy for both linear and nonlinear components.

4.4 GENERAL MPFPP EQUATIONS

Below are included the generic equation which characterize all models. Specific instances are discussed in section 4.2. Their source code can be found at the following URLs:

http://cns.bu.edu/~ezaydens/data/soft/mpfpp_m1.zip

http://cns.bu.edu/~ezaydens/data/soft/mpfpp_m2hh.zip

Marginalized Particle Filter with Point Process observations system

$$\begin{aligned}
x_{k+1}^n &= f_k^n(x_k^n) + A_k^n(x_k^n) \cdot x_k^l + G_k^n(x_k^n) \varepsilon_k^n \\
x_{k+1}^l &= f_k^l(x_k^n) + A_k^l(x_k^n) \cdot x_k^l + G_k^l(x_k^n) \varepsilon_k^l \\
y_k &= \text{Pois}(\lambda_k \Delta) \\
\lambda_k &= \exp\{h_k(x_k^n) + C_k(x_k^n) \cdot x_k^l\} \\
\begin{bmatrix} \varepsilon_k^n \\ \varepsilon_k^l \end{bmatrix} &\sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} Q_k^n & Q_k^{\text{ln}} \\ (Q_k^{\text{ln}})^T & Q_k^l \end{bmatrix}\right)
\end{aligned}$$

Nonlinear estimation equations

One-step prediction:

$$x_{k+1}^n = f_k^n(x_k^n) + A_k^n(x_k^n) \cdot x_k^l + G_k^n(x_k^n) \varepsilon_k^n.$$

Weight update:

$$\begin{aligned}
w_k &= w_{k-1} \cdot p(y_k | x_{0:k}^n) \\
p(y_k | x_{0:k}^n) &= p(\Delta N_k | x_{0:k}^n) = (\lambda_k^{(n)} \cdot \Delta)^{\Delta N_k} \cdot e^{-\lambda_k^{(n)} \cdot \Delta}, \\
\lambda_k^{(n)} &= \exp\left\{h_k(x_k^n) + C_k \cdot x_k^l + \frac{1}{2} C_k \cdot P_{k|k-1} \cdot C_k^T\right\}.
\end{aligned}$$

Estimate computation:

$$\hat{x}_k^n = \sum_{i=1}^M w_k^{(i)} x_k^{n,(i)}.$$

Linear estimation equations

One-step prediction:

$$\begin{aligned}
\hat{x}_{k+1|k}^{l,(i)} &= \bar{A}_k^{l,(i)} \hat{x}_{k|k}^{l,(i)} + G_k^{l,(i)} (Q_k^{\text{ln}})^T (G_k^{n,(i)} Q_k^n)^{-1} z_k^{(i)} + f_k^{l,(i)} + L_k^{(i)} (z_k^{(i)} - A_k^{n,(i)} \hat{x}_{k|k}^{l,(i)}), \\
P_{k+1|k}^{(i)} &= \bar{A}_k^{l,(i)} P_{k|k}^{(i)} (\bar{A}_k^{l,(i)})^T + G_k^{l,(i)} \bar{Q}^l (G_k^{l,(i)})^T - L_k^{(i)} N_k^{(i)} (L_k^{(i)})^T, \\
N_k^{(i)} &= A_k^{n,(i)} P_{k|k}^{(i)} (A_k^{n,(i)})^T + G_k^{n,(i)} Q_k^n (G_k^{n,(i)})^T, \\
L_k^{(i)} &= \bar{A}_k^{l,(i)} P_{k|k}^{(i)} (A_k^{n,(i)})^T (N_k^{(i)})^{-1}, \\
z_k^{(i)} &= x_{k+1}^n - f_k^{n,(i)}, \\
\bar{A}_k^{l,(i)} &= A_k^{l,(i)} - G_k^{l,(i)} (Q_k^{\text{ln}})^T (G_k^{n,(i)} Q_k^n)^{-1} A_k^{n,(i)}, \\
\bar{Q}^l &= Q_k^l - (Q_k^{\text{ln}})^T (Q_k^n)^{-1} Q_k^{\text{ln}}.
\end{aligned}$$

Posterior update:

$$\hat{\mathbf{x}}_{k|k}^{l,(i)} = \hat{\mathbf{x}}_{k|k-1}^{l,(i)} + \mathbf{P}_{k|k}^{l,(i)} \left[\frac{\partial \log \lambda_k}{\partial \mathbf{x}_k^l} (\Delta N_k - \lambda_k \cdot \Delta) \right]_{\hat{\mathbf{x}}_{k|k-1}^{l,(i)}},$$

$$\left(\mathbf{P}_{k|k}^{l,(i)} \right)^{-1} = \left(\mathbf{P}_{k|k-1}^{l,(i)} \right)^{-1} + \left[\left(\frac{\partial \log \lambda_k}{\partial \mathbf{x}_k^l} \right)^T (\lambda_k \cdot \Delta) \left(\frac{\partial \log \lambda_k}{\partial \mathbf{x}_k^l} \right) - (\Delta N_k - \lambda_k \cdot \Delta) \frac{\partial^2 \log \lambda_k}{\partial \mathbf{x}_k^l \partial (\mathbf{x}_k^l)^T} \right]_{\hat{\mathbf{x}}_{k|k-1}^{l,(i)}}.$$

Estimate computation:

$$\hat{\mathbf{x}}_{k|k}^l = \sum_{i=1}^M w_k^{(i)} \hat{\mathbf{x}}_{k|k}^{l,(i)},$$

$$\hat{\mathbf{P}}_{k|k} = \sum_{i=1}^M w_k^{(i)} \left(\hat{\mathbf{P}}_{k|k}^{(i)} + (\hat{\mathbf{x}}_{k|k}^{l,(i)} - \hat{\mathbf{x}}_{k|k}^l)(\hat{\mathbf{x}}_{k|k}^{l,(i)} - \hat{\mathbf{x}}_{k|k}^l)^T \right).$$

5 CONCLUSIONS

This thesis developed point process methodologies for human peripheral nervous system research and applied them to the recordings of muscle sympathetic nerve activity. The point process paradigm operates using the stereotypical neuronal electrical events called action potentials or spikes. The thick tungsten microelectrodes that percutaneously record human peripheral nerve activity produce signals that are corrupted with high levels of noise, and the identification of individual action potentials buried in these noisy signals has only recently become available.

The advent of the new spike extraction techniques provided answers to some of the longstanding peripheral nervous system researchers questions. For example, traditional analysis showed that muscle sympathetic nerve activity contains pulse-synchronous bursts, but the temporal smoothing inherent in the analysis procedure prevented the internal structure of these bursts from being examined. Once spikes became available, their distributions were computed, the individual action potential shapes were classified, and a specific recruitment order of the individual muscle sympathetic nerve fibers was suggested.

Other research problems remained unaddressed. For example, the sympathetic activity occurring between bursts has not been assessed. The physiological mechanisms guiding the recruitment of fibers, the timing of spikes, and the overall regulation of the burst structure are still unknown. To answer these questions, a thorough analysis of both the structure of the peripheral nervous system spike trains and their relationship to the hemodynamic variables is necessary.

Point process modeling and analysis framework handles both. Structurally, very few limitations exist on the types of spike trains point process models can handle; models were successfully constructed for very sparse spiking as well as for the very high-frequency firing. This is important for analyzing the peripheral nervous system data because individual fibers can fire as rarely as once in several seconds, but the multifiber firing rates can reach into hundreds of spikes/sec. In fact, the model we construct in Chapter 2 utilizes this versatility of the point process methodology to relate the fast to the slow spiking within a single sympathetic recording.

Point process framework also offers a structured and statistically rigorous approach to describing the relationship between nerve activity and the putative covariates. Both explanatory and predictive models can be constructed. The former quantifies the information that spike trains contain about the covariates, while the latter predicts covariate values directly from the spike trains. Ability to construct both types of models is crucial for the peripheral nervous system analysis because many physiological functions involve feedback. For instance, a baroreflex negative feedback loop is vital to the ability of the muscle sympathetic nerve activity to maintain cardiovascular homeostasis. Moreover, point processes framework allows dynamic estimation of covariates from the observed spike trains; an important faculty in modeling transient changes to spiking-covariate relationships induced by some experimental conditions

5.1 Review of Thesis Results

Chapter 1 showed that previous research into characterizing neural spiking activity using point processes was highly fruitful. The work presented in this thesis

expands the types of neural systems that point process modeling and estimation framework is applied to and improves the theoretical component of the framework to handle the complexities encountered in the human peripheral nervous system data.

In Chapters 2 and 3 we successfully characterized and modeled the high-frequency spike trains representing the muscle sympathetic nerve activity (MSNA). Point process analysis framework allowed us to approach it in two steps. First, we focused on the local variability of the spike trains, describing the distribution of spikes within individual cardiac cycles and modeling this distribution as a function of the low-frequency variability in the spike trains. In the second step we characterized the sympathetic spike train variability across multiple frequency bands using three broad classes of covariates: recording duration, instantaneous phases of hemodynamic variables, and cardiac cycle phase. We showed that the sympathetic spike train variability is very high: we only were able to capture about 10% of it on average despite our covariates spanning the entire physiologically relevant frequency spectrum.

We also developed an algorithm that uses random samples to estimate the trajectory of a random process from the spike train output the process has produced. The algorithm leveraged the known substructure within the process model in order to break its posterior probability into a linear and nonlinear components. Those were then estimated separately; the methods best suited for each one were used, increasing accuracy and reducing computational load. We used the new algorithm to concurrently estimate both the membrane dynamics and the maximal conductance parameters for a simulated Hodgkin-Huxley biophysical model.

5.2 *Future Research Directions*

The signal used to analyze human peripheral nerve activity has traditionally had low temporal resolution. Therefore, access to the millisecond-scale information contained in the spike trains created abundant opportunities for exciting new work. First, some of the questions addressed with the traditional methods could now be studied using spike trains. For instance, the description of the multifiber MSNA bursts that consisted merely of start and peak times as well as height and total area can now be substantially expanded to include multifiber spike distributions as well as individual fibers' action potential placement. Additionally, work on examining and characterizing the sizes, types, and activation patterns of individual nerves can also now begin. Some early steps have already been taken by Salmanpour, Steinback, and others [122, 110, 111, 108], but the identification of physiological mechanisms that turn on and regulate different neuron types across experimental conditions remains an important unsolved problem.

Point process modeling framework is well-suited to address these challenges. Methods for recording the activity of many neurons simultaneously have been developed [30]. To analyze these data, techniques were devised that modeled the individual neurons' probability of spike emission as a function of the spiking history by the entire ensemble of neurons [126]. Recently, point process models were extended to study functional connectivity within neuronal networks and applied to the recording of rat hippocampal place cells [95]. The ability to identify cells that communicate with each other solely from the cells' output is invaluable to the human peripheral nervous system researchers because they cannot record directly from the individual nerve fibers.

Applying these techniques to the newly identified sympathetic fiber types will deepen our understanding of interactivity patterns and suggest physiologies responsible for this function.

Instead of predicting the spikes, we can use individual neuron's history of firing, the activity of neuronal ensembles, and other stimulus signals to predict the subjects' behavioral patterns and physiological processes. This method, called neural decoding, is ubiquitous in computational neuroscience, most notably in brain-machine interface and prosthetics research, where decoders control a wide range of peripheral devices. Point process models were also successful in this setting, using spike trains recorded from the sensory, pre-motor, and motor cortices to predict the position, velocity, and acceleration trajectories of various robotic actuators. Decoding can serve a multitude of useful functions in the peripheral nervous system research. For example, vascular transduction properties of the muscle sympathetic nerves can be studied by decoding their effective control of the arterial blood pressure. The response properties of sensory nerves can be researched by predicting the stimulus intensity from, for example, the skin sympathetic nerve activity.

Point processes have also been used to perform decoding dynamically, estimating for example, the internal currents of an individual neuron from the spikes the neuron produced. Access to the peripheral nervous system spike trains now allows direct application of this technique to studying the physiology of individual nerve types within the multifiber nerve bundles. Constructing accurate models of cross-membrane currents will help link the known anatomy of the human peripheral nerves (sizes, capacitances,

etc.) to their physiology and function. These can then be compared to similar studies in other mammals, illuminating the *in vivo* functionality of various nerve branches. Marginalized particle filters with point process observations that we developed will be helpful in these studies as they can reduce the computational complexity of some estimation problems as well as increase the accuracy of the estimates when compared with other techniques.

The marginalized particle filter with point process observations can be further customized to fit the properties of the signals that occur in peripheral nervous system research. In this thesis we separated the conditionally linear and the marginally nonlinear posterior densities, applying a Gaussian approximation to the former and using particles to approximate the latter. Other ways of processing linear and nonlinear components can be envisioned. For instance, the nonlinear observation model can be linearized using extended or unscented Kalman filter. Moreover, if the nonlinearities in the state transition functions are not severe, one can linearize those as well and eliminate the random samples altogether. This combination of filters would significantly increase the computational speed, but is likely to reduce accuracy. Additionally, point process analogues of the recursive least squares and of the steepest descent filters [56] can be used for systems with favorable properties.

In this thesis we used marginalization to help us estimate the state variables from past and present spike train observations. However, marginalizing the filtering distribution is just one of the many available modifications that the state-space paradigm for point process observations allows. For example, future spike train observations can

be used to compute the marginalized smoothing distribution and to improve estimation accuracy even further. Additionally, a wide array of other signal processing techniques can be incorporated into the point process dynamical estimation schemes either together with the marginalization or separate from it. Jump Markov linear systems [33], dependent noise processes [104], and fine-tuning the importance distribution [116] are just a few that come to mind.

BIBLIOGRAPHY

- [1] E. D. Adrian, D. W. Bronk, and Gilbert Phillips. Discharges in mammalian sympathetic nerves. *The Journal of physiology*, 74(2):115–133, February 1932.
- [2] D. L. Alspach and H. W. Sorenson. Nonlinear Bayesian estimation using Gaussian sum approximations. *IEEE Transactions on Automatic Control*, 17(4):439–448, August 1972.
- [3] P. O. Andersson. Comparative vascular effects of stimulation continuously and in bursts of the sympathetic nerves to cat skeletal muscle. *Acta physiologica Scandinavica*, 118(4):343–348, August 1983.
- [4] S. Ando, T. Imaizumi, and A. Takeshita. Effects of patterns of sympathetic nerve stimulation on vasoconstricting responses in the hindquarter of rabbits. *Journal of the autonomic nervous system*, 45(3):225–233, December 1993.
- [5] C. Andrieu and S. J. Godsill. A particle filter for model based audio source separation. In *International Workshop on Independent Component Analysis and Blind Signal Separation (ICA 2000)*, pages 381–386, 2000.
- [6] Christophe Andrieu and Arnaud Doucet. Particle filtering for partially observed Gaussian state space models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(4):827–836, October 2002.
- [7] S. N. Baker, J. M. Kilner, E. M. Pinches, and R. N. Lemon. The role of synchrony and oscillations in the motor output. *Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale*, 128(1-2):109–117, September 1999.
- [8] Robert M. Berne and Matthew N. Levy. *Cardiovascular Physiology*. Mosby Physiology Monograph Series, St. Louis, MO, eighth edition, 2001.
- [9] A. Boczek-Funcke, H. J. Häbler, W. Jänig, and M. Michaelis. Respiratory modulation of the activity in sympathetic neurones supplying muscle, skin and pelvic organs in the cat. *The Journal of physiology*, 449:333–361, April 1992.
- [10] A. Borovik, V. Golubinskaya, O. Tarasova, C. Aalkjaer, and H. Nilsson. Phase resetting of arterial vasomotion by burst stimulation of perivascular nerves. *Journal of vascular research*, 42(2):165–173, 2005.
- [11] G. Box, G. Jenkins, and G. Reinsel. *Time Series Analysis: Forecasting and Control*. Prentice Hall, third edition, 1994.
- [12] Toni Breskovic, Craig D. Steinback, Aryan Salmanpour, J. Kevin Shoemaker, and Zeljko Dujic. Recruitment pattern of sympathetic neurons during breath-holding at

- different lung volumes in apnea divers and controls. *Autonomic neuroscience : basic & clinical*, 164(1-2):74–81, October 2011.
- [13] A. E. Brockwell, A. L. Rojas, and R. E. Kass. Recursive Bayesian Decoding of Motor Cortical Signals by Particle Filtering. *Journal of Neurophysiology*, 91(4):1899–1907, April 2004.
- [14] D. W. Bronk, L. K. Ferguson, R. Margaria, and D. Y. Solandt. The activity of the cardiac sympathetic centers. *American Journal of Physiology – Legacy Content*, 117(2):237–249, September 1936.
- [15] E. N. Brown, R. Barbieri, U. T. Eden, and L. M. Frank. *Likelihood methods for neural spike train data analysis*, chapter 9, pages 253–286. CRC Press, London, 2003.
- [16] R. J. Brychta, W. Charoensuk, L. Bernardi, R. Furlan, R. Shiavi, and A. Diedrich. Spectral analysis of multiunit action potential trains of muscle sympathetic nerve activity in humans. In *Computers in Cardiology, 2002*, pages 457–460. IEEE, 2002.
- [17] Robert J. Brychta, Richard Shiavi, David Robertson, and André Diedrich. Spike detection in human muscle sympathetic nerve activity using the kurtosis of stationary wavelet transform coefficients. *Journal of neuroscience methods*, 160(2):359–367, March 2007.
- [18] Olivier Cappé, Simon J. Godsill, and Eric Moulines. An Overview of Existing Methods and Recent Advances in Sequential Monte Carlo. *Proceedings of the IEEE*, 95(5):899–924, May 2007.
- [19] J. T. Carlson, J. Hedner, M. Elam, H. Ejnell, J. Sellgren, and B. G. Wallin. Augmented resting sympathetic activity in awake patients with obstructive sleep apnea. *Chest*, 103(6):1763–1768, June 1993.
- [20] C. E. Carr and M. Konishi. A circuit for detection of interaural time differences in the brain stem of the barn owl. *The Journal of Neuroscience*, 10(10):3227–3246, October 1990.
- [21] George Casella and Christian P. Robert. Rao-Blackwellisation of sampling schemes. *Biometrika*, 83(1):81–94, March 1996.
- [22] Rong Chen, Xiaodong Wang, and J. S. Liu. Adaptive joint detection and decoding in flat-fading channels via mixture Kalman filtering. *IEEE Transactions on Information Theory*, 46(6):2079–2094, September 2000.

- [23] H. P. Clamann and E. Henneman. Electrical measurement of axon diameter and its use in relating motoneuron size to critical firing level. *Journal of neurophysiology*, 39(4):844–851, July 1976.
- [24] Jacob Cohen, Patricia Cohen, Stephen G. West, and Leona S. Aiken. *Applied Multiple Regression/Correlation Analysis for the Behavioral Sciences, 3rd Edition*. Routledge, third edition, August 2002.
- [25] W. Delius, K. E. Hagbarth, A. Hongell, and B. G. Wallin. General characteristics of sympathetic activity in human muscle nerves. *Acta physiologica Scandinavica*, 84(1):65–81, January 1972.
- [26] W. Delius, K. E. Hagbarth, A. Hongell, and B. G. Wallin. Manoeuvres affecting sympathetic outflow in human muscle nerves. *Acta physiologica Scandinavica*, 84(1):82–94, January 1972.
- [27] Jerome A. Dempsey, A. William Sheel, Claudette M. St Croix, and Barbara J. Morgan. Respiratory influences on sympathetic vasomotor outflow in humans. *Respiratory physiology & neurobiology*, 130(1):3–20, March 2002.
- [28] Luc Devroye. *Non-Uniform Random Variate Generation*. Springer New York, 1986.
- [29] André Diedrich, Warakorn Charoensuk, Robert J. Brychta, Andrew C. Ertl, and Richard Shiavi. Analysis of raw microneurographic recordings based on wavelet denoising technique and classification algorithm: wavelet analysis in microneurography. *IEEE transactions on bio-medical engineering*, 50(1):41–50, January 2003.
- [30] J. Donoghue. Bridging the Brain to the World: A Perspective on Neural Interface Systems. *Neuron*, 60(3):511–521, November 2008.
- [31] Arnaud Doucet, Nando Freitas, and Neil J. Gordon. An Introduction to Sequential Monte Carlo Methods. In Arnaud Doucet, Nando Freitas, and Neil Gordon, editors, *Sequential Monte Carlo Methods in Practice*, Statistics for Engineering and Information Science, pages 3–14. Springer New York, 2001.
- [32] Arnaud Doucet, Simon J. Godsill, and Christophe Andrieu. On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, 10(3):197–208, July 2000.
- [33] Arnaud Doucet, N. J. Gordon, and V. Krishnamurthy. Particle filters for state estimation of jump Markov linear systems. *IEEE Transactions on Signal Processing*, 49(3):613–624, March 2001.

- [34] Arnaud Doucet and Adam M. Johansen. A tutorial on particle filtering and smoothing: fifteen years later. In *Oxford Handbook of Nonlinear Filtering*, pages 656–704, 2008.
- [35] Uri T. Eden, Loren M. Frank, Riccardo Barbieri, Victor Solo, and Emery N. Brown. Dynamic Analysis of Neural Encoding by Point Process Adaptive Filtering. *Neural Computation*, 16(5):971–998, May 2004.
- [36] M. Elam, Y. B. Sverrisdottir, B. Rundqvist, D. McKenzie, B. Gunnar Wallin, and V. G. Macefield. Pathological sympathoexcitation: how is it achieved? *Acta physiologica Scandinavica*, 177(3):405–411, March 2003.
- [37] Ayla Ergün, Riccardo Barbieri, Uri T. Eden, Matthew A. Wilson, and Emery N. Brown. Construction of point process adaptive filter algorithms for neural systems using sequential Monte Carlo methods. *IEEE transactions on bio-medical engineering*, 54(3):419–428, March 2007.
- [38] M. Esler. The sympathetic system and hypertension. *American Journal of Hypertension*, 13(6):S99–S105, June 2000.
- [39] J. Fagius. Sympathetic nerve activity in metabolic control “some basic concepts. *Acta Physiologica Scandinavica*, 177(3):337–343, 2003.
- [40] J. Fagius, G. Sundlöf, and B. G. Wallin. Variation of sympathetic reflex latency in man. *Journal of the autonomic nervous system*, 21(2-3):157–165, December 1987.
- [41] J. Fagius and B. G. Wallin. Sympathetic reflex latencies and conduction velocities in normal man. *Journal of the neurological sciences*, 47(3):433–448, September 1980.
- [42] J. S. Floras. Sympathetic activation in human heart failure: diverse mechanisms, therapeutic opportunities. *Acta Physiologica Scandinavica*, 177(3):391–398, 2003.
- [43] R. Furlan, A. Porta, F. Costa, J. Tank, L. Baker, R. Schiavi, D. Robertson, A. Malliani, and R. Mosqueda-Garcia. Oscillatory patterns in sympathetic neural discharge and cardiovascular variables during orthostatic stimulus. *Circulation*, 101(8):886–892, February 2000.
- [44] A. Gelb, editor. *Applied Optimal Estimation*. MIT Press, Boston, MA, 1974.
- [45] Simon J. Godsill, Arnaud Doucet, and Mike West. Monte Carlo Smoothing for Nonlinear Time Series. *Journal of the American Statistical Association*, 99(465):438–449, 2004.

- [46] N. J. Gordon, D. J. Salmond, and A. F. M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F (Radar and Signal Processing)*, 140(2):107–113, April 1993.
- [47] G. Grassi and M. Esler. How to assess sympathetic activity in humans. *Journal of hypertension*, 17(6):719–734, June 1999.
- [48] J. P. Greenwood, J. B. Stoker, and D. A. Mary. Single-unit sympathetic discharge : quantitative assessment in human hypertensive disease. *Circulation*, 100(12):1305–1310, September 1999.
- [49] F. Gustafsson, F. Gunnarsson, Niclas Bergman, U. Forssell, J. Jansson, R. Karlsson, and P. J. Nordlund. Particle filters for positioning, navigation, and tracking. *IEEE Transactions on Signal Processing*, 50(2):425–437, February 2002.
- [50] Fredrik Gustafsson and Thomas Schön. Particle Filters For System Identification Of State-Space Models Linear In Either Parameters Or States. In *13th International Federation of Automatic Control Symposium on System Identification*, pages 1287–1292, 2003.
- [51] H. J. Häbler, W. Jänig, M. Krummel, and O. A. Peters. Reflex patterns in postganglionic neurons supplying skin and skeletal muscle of the rat hindlimb. *Journal of neurophysiology*, 72(5):2222–2236, November 1994.
- [52] H. J. Häbler, W. Jänig, and M. Michaelis. Respiratory modulation in the activity of sympathetic neurones. *Progress in neurobiology*, 43(6):567–606, August 1994.
- [53] K. E. Hagbarth and A. B. Vallbo. Mechanoreceptor activity recorded percutaneously with semi-microelectrodes in human peripheral nerves. *Acta physiologica Scandinavica*, 69(1):121–122, 1967.
- [54] J. E. Handschin. Monte Carlo techniques for prediction and filtering of non-linear stochastic processes. *Automatica*, 6(4):555–563, July 1970.
- [55] J. E. Handschin and D. Q. Mayne. Monte Carlo techniques to estimate the conditional expectation in multi-stage non-linear filtering. *International Journal of Control*, 9(5):547–559, May 1969.
- [56] Simon O. Haykin. *Adaptive Filter Theory (4th Edition)*. Prentice Hall, 4 edition, September 2001.
- [57] A. Hodgkin and A. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *Bulletin of Mathematical Biology*, 52(1):25–71, January 1952.

- [58] G. Horeysek and W. Jänig. Reflex activity in postganglionic fibres within skin and muscle nerves elicited by somatic stimuli in chronic spinal cats. *Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale*, 21(2):155–168, 1974.
- [59] Masashi Ichinose, Mitsuru Saito, Narihiko Kondo, and Takeshi Nishiyasu. Time-dependent modulation of arterial baroreflex control of muscle sympathetic nerve activity during isometric exercise in humans. *American journal of physiology. Heart and circulatory physiology*, 290(4), April 2006.
- [60] T. Ikeda, S. Iwase, M. Saito, and T. Mano. Effects of positive and negative pressure breathing on muscle sympathetic nerve activity in humans. *Aviation, space, and environmental medicine*, 68(6):494–498, June 1997.
- [61] Y. Ikeda, T. Kawada, M. Sugimachi, O. Kawaguchi, T. Shishido, T. Sato, H. Miyano, W. Matsuura, J. Alexander, and K. Sunagawa. Neural arc of baroreflex optimizes dynamic pressure regulation in achieving both stability and quickness. *The American journal of physiology*, 271(3 Pt 2), September 1996.
- [62] K. Ito and Kaiqi Xiong. Gaussian filters for nonlinear filtering problems. *IEEE Transactions on Automatic Control*, 45(5):910–927, May 2000.
- [63] E. M. Izhikevich. Simple model of spiking neurons. *IEEE Transactions on Neural Networks*, 14(6):1569–1572, November 2003.
- [64] W. Jänig. Organization of the lumbar sympathetic outflow to skeletal muscle and skin of the cat hindlimb and tail. *Reviews of physiology, biochemistry and pharmacology*, 102:119–213, 1985.
- [65] Andrew H. Jazwinski. *Stochastic Processes and Filtering Theory*. Academic Press, 111 Fifth Ave., New York, NY 10003, January 1970.
- [66] C. B. Jenq, L. L. Jenq, and R. E. Coggeshall. Numerical patterns of axon regeneration that follow sciatic nerve crush in the neonatal rat. *Experimental neurology*, 95(2):492–499, February 1987.
- [67] Claude Julien. The enigma of Mayer waves: Facts and models. *Cardiovascular research*, 70(1):12–21, April 2006.
- [68] Simon J. Julier and Jeffrey K. Uhlmann. *A New Extension of the Kalman Filter to Nonlinear Systems*, pages 182–193. July 1997.
- [69] R. E. Kalman. A New Approach to Linear Filtering and Prediction Problems. *Transactions of the American Society of Mechanical Engineers, Journal of Basic Engineering*, 82(Series D):35–45, 1960.

- [70] R. E. Kalman and R. S. Bucy. New results in linear filtering and prediction theory. In *Transactions of the American Society of Mechanical Engineers, Journal of Basic Engineering*, 1961.
- [71] R. J. Kaminski, G. A. Meyer, and D. L. Winter. Sympathetic unit activity associated with Mayer waves in the spinal dog. *The American journal of physiology*, 219(6):1768–1771, December 1970.
- [72] A. Kamiya, D. Michikami, Q. Fu, Y. Niimi, S. Iwase, T. Mano, and A. Suzumura. Static handgrip exercise modifies arterial baroreflex control of vascular sympathetic outflow in humans. *American journal of physiology. Regulatory, integrative and comparative physiology*, 281(4), October 2001.
- [73] Atsunori Kamiya, Junichiro Hayano, Toru Kawada, Daisaku Michikami, Kenta Yamamoto, Hideto Ariumi, Syuji Shimizu, Kazunori Uemura, Tadayoshi Miyamoto, Takeshi Aiba, Kenji Sunagawa, and Masaru Sugimachi. Low-frequency oscillation of sympathetic nerve activity decreases during development of tilt-induced syncope preceding sympathetic withdrawal and bradycardia. *American journal of physiology. Heart and circulatory physiology*, 289(4), October 2005.
- [74] R. Karlsson, F. Gusafsson, and T. Karlsson. Particle filtering and Cramer-Rao lower bound for underwater navigation. In *IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003.*, volume 6, pages VI–65–8 vol.6. IEEE, April 2003.
- [75] R. Karlsson, T. Schön, and F. Gustafsson. Complexity analysis of the marginalized particle filter. *IEEE Transactions on Signal Processing*, 53(11):4408–4411, November 2005.
- [76] Marc P. Kaufman and Hubert V. Forster. *Reflexes Controlling Circulatory, Ventilatory and Airway Responses to Exercise*, chapter 12, pages 381–447. oxford University Press for the American Physiological Society, New York, 1996.
- [77] Genshiro Kitagawa. A self-organizing state-space model. *Journal of the American Statistical Association*, 93(443):1203–1215, September 1998.
- [78] W. H. Levison, G. O. Barnett, and W. D. Jackson. Nonlinear analysis of the baroreceptor reflex system. *Circulation research*, 18(6):673–682, June 1966.
- [79] Ping Li, Roger Goodall, and Visakan Kadiramanathan. Parameter estimation of railway vehicle dynamic model using Rao-Blackwellised particle filter. In *European Control Conference*, pages 1–4, 2003.
- [80] Jane Liu and Mike West. Combined Parameter and State Estimation in Simulation-Based Filtering. In Arnaud Doucet, Nando Freitas, and Neil Gordon,

- editors, *Sequential Monte Carlo Methods in Practice*, Statistics for Engineering and Information Science, pages 197–223. Springer New York, 2001.
- [81] V. G. Macefield, M. Elam, and B. G. Wallin. Firing properties of single postganglionic sympathetic neurones recorded in awake human subjects. *Autonomic neuroscience : basic & clinical*, 95(1-2):146–159, January 2002.
- [82] V. G. Macefield, B. G. Wallin, and A. B. Vallbo. The discharge behaviour of single vasoconstrictor motoneurons in human muscle nerves. *The Journal of Physiology*, 481 (Pt 3):799–809, December 1994.
- [83] Vaughan G. Macefield and Gunnar B. Wallin. Firing properties of single vasoconstrictor neurones in human subjects with high levels of muscle sympathetic activity. *J Physiol*, 516(1):293–301, April 1999.
- [84] A. Malliani, M. Pagani, F. Lombardi, and S. Cerutti. Cardiovascular neural regulation explored in the frequency domain. *Circulation*, 84(2):482–492, August 1991.
- [85] Tadaaki Mano, Satoshi Iwase, and Shinobu Toma. Microneurography as a tool in clinical neurophysiology to investigate peripheral neural traffic in humans. *Clinical neurophysiology*, 117(11):2357–2384, November 2006.
- [86] Petra Zubin Z. Maslov, Toni Breskovic, Danielle N. Brewer, J. Kevin Shoemaker, and Zeljko Dujic. Recruitment pattern of sympathetic muscle neurons during premature ventricular contractions in heart failure patients and controls. *American journal of physiology. Regulatory, integrative and comparative physiology*, 303(11), December 2012.
- [87] R. M. McAllen and K. M. Spyer. The baroreceptor input to cardiac vagal motoneurons. *The Journal of physiology*, 282:365–374, September 1978.
- [88] D. I. McCloskey and E. K. Potter. Excitation and inhibition of cardiac vagal motoneurons by electrical stimulation of the carotid sinus nerve. *The Journal of physiology*, 316:163–175, July 1981.
- [89] Liang Meng, Mark A. Kramer, and Uri T. Eden. A sequential Monte Carlo approach to estimate biophysical neural models from spikes. *Journal of Neural Engineering*, 8(6):065006+, December 2011.
- [90] N. Montano, C. Cogliati, V. J. da Silva, T. Gnecci-Ruscione, M. Massimini, A. Porta, and A. Malliani. Effects of spinal section and of positive-feedback excitatory reflex on sympathetic and heart rate variability. *Hypertension*, 36(6):1029–1034, December 2000.

- [91] N. Montano, T. Gneccchi-Ruscione, A. Porta, F. Lombardi, A. Malliani, and S. M. Barman. Presence of vasomotor and respiratory rhythms in the discharge of single medullary neurons involved in the regulation of cardiovascular system. *Journal of the autonomic nervous system*, 57(1-2):116–122, February 1996.
- [92] Hisayoshi Murai, Shigeo Takata, Michiro Maruyama, Manabu Nakano, Daisuke Kobayashi, Kan-ichi Otowa, Masayuki Takamura, Toyoshi Yuasa, Satoru Sakagami, and Shuichi Kaneko. The activity of a single muscle sympathetic vasoconstrictor nerve unit is affected by physiological stress in humans. *American Journal of Physiology, Heart and Circulatory Physiology*, 290(2):H853–860, February 2006.
- [93] J. Nagumo, S. Arimoto, and S. Yoshizawa. An Active Pulse Transmission Line Simulating Nerve Axon. *Proceedings of the Institute of Radio Engineers*, 50(10):2061–2070, October 1962.
- [94] Per-Johan Nordlund. *Sequential Monte Carlo Filters and Integrated Navigation*. PhD thesis, Division of Automatic Control, Department of Electrical Engineering, Linköpings universitet, 2002.
- [95] Murat Okatan, Matthew A. Wilson, and Emery N. Brown. Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural computation*, 17(9):1927–1961, September 2005.
- [96] M. Pagani, N. Montano, A. Porta, A. Malliani, F. M. Abboud, C. Birkett, and V. K. Somers. Relationship between spectral components of cardiovascular variabilities and direct measures of muscle sympathetic nerve activity in humans. *Circulation*, 95(6):1441–1448, March 1997.
- [97] Michael K. Pitt and Neil Shephard. Filtering via Simulation: Auxiliary Particle Filters. *Journal of the American Statistical Association*, 94(446):590–599, June 1999.
- [98] B. Povlsen, N. Stankovic, P. Danielsson, and C. Hildebrand. Fiber composition of the lateral plantar and superficial peroneal nerves in the rat foot. *Anatomy and embryology*, 189(5):393–399, May 1994.
- [99] G. Preiss and C. Polosa. Patterns of sympathetic neuron activity associated with Mayer waves. *The American journal of physiology*, 226(3):724–730, March 1974.
- [100] M. Rosenbaum and D. Race. Frequency-response characteristics of vascular resistance vessels. *The American journal of physiology*, 215(6):1397–1402, December 1968.
- [101] Donald B. Rubin. A noniterative sampling/importance resampling alternative to the data augmentation algorithm for creating a few imputations when the fraction of

- missing information is modest: The SIR algorithm. *Journal of American Statistical Association*, 82:543–546, 1987.
- [102] Donald B. Rubin. *Using the SIR algorithm to simulate posterior distribution*, pages 395–402. Oxford University Press, Oxford, U.K., 1988.
- [103] Kathy L. Ryan, Caroline A. Rickards, Carmen Hinojosa-Laborde, William H. Cooke, and Victor A. Convertino. Arterial pressure oscillations are not associated with muscle sympathetic nerve activity in individuals exposed to central hypovolaemia. *The Journal of physiology*, 589(Pt 21):5311–5322, November 2011.
- [104] S. Saha and F. Gustafsson. Marginalized particle filter for dependent Gaussian noise processes. In *2012 IEEE Aerospace Conference*, pages 1–6. IEEE, March 2012.
- [105] M. Saito, N. Foldager, T. Mano, S. Iwase, Y. Sugiyama, and M. Oshima. Sympathetic control of hemodynamics during moderate head-up tilt in human subjects. *Environmental medicine : annual report of the Research Institute of Environmental Medicine, Nagoya University*, 41(2):151–155, December 1997.
- [106] M. Saito, T. Mano, H. Abe, and S. Iwase. Responses in muscle sympathetic nerve activity to sustained hand-grips of different tensions in humans. *European journal of applied physiology and occupational physiology*, 55(5):493–498, 1986.
- [107] M. Saito, M. Naito, and T. Mano. Different responses in skin and muscle sympathetic nerve activity to static muscle contraction. *Journal of applied physiology (Bethesda, Md. : 1985)*, 69(6):2085–2090, December 1990.
- [108] A. Salmanpour, M. F. Frances, R. Goswami, and J. K. Shoemaker. Sympathetic neural recruitment patterns during the Valsalva maneuver. In *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 6951–6954. IEEE, August 2011.
- [109] Aryan Salmanpour, Lyndon J. Brown, and Shoemaker. Spike detection in human muscle sympathetic nerve activity using a matched wavelet approach. *Journal of Neuroscience Methods*, 193(2):343–355, November 2010.
- [110] Aryan Salmanpour, Lyndon J. Brown, Craig D. Steinback, Charlotte W. Usselman, Ruma Goswami, and J. Kevin Shoemaker. Relationship between size and latency of action potentials in human muscle sympathetic nerve activity. *Journal of neurophysiology*, 105(6):2830–2842, June 2011.
- [111] Aryan Salmanpour and J. Kevin Shoemaker. Baroreflex mechanisms regulating the occurrence of neural spikes in human muscle sympathetic nerve activity. *Journal of Neurophysiology*, 107(12):3409–3416, June 2012.

- [112] Takayuki Sato, Toru Kawada, Masashi Inagaki, Toshiaki Shishido, Masaru Sugimachi, and Kenji Sunagawa. Dynamics of sympathetic baroreflex control of arterial pressure in rats. *American journal of physiology. Regulatory, integrative and comparative physiology*, 285(1), July 2003.
- [113] J. P. Saul, R. F. Rea, D. L. Eckberg, R. D. Berger, and R. J. Cohen. Heart rate and muscle sympathetic nerve variability during reflex changes of autonomic activity. *The American journal of physiology*, 258(3 Pt 2), March 1990.
- [114] A. M. SCHER and A. C. YOUNG. Servoanalysis of carotid sinus reflex effects on peripheral resistance. *Circulation research*, 12:152–162, February 1963.
- [115] H. Schmalbruch. Fiber composition of the rat sciatic nerve. *The Anatomical record*, 215(1):71–81, May 1986.
- [116] T. B. Schön, F. Gustafsson, and P. J. Nordlund. Marginalized particle filters for mixed linear/nonlinear state-space models. *IEEE Transactions on Signal Processing*, 53(7):2279–2289, July 2005.
- [117] D. R. Seals. Sympathetic neural discharge and vascular resistance during exercise in humans. *Journal of applied physiology (Bethesda, Md. : 1985)*, 66(5):2472–2478, May 1989.
- [118] J. P. Segundo, G. P. Moore, L. J. Stensaas, and T. H. Bullock. Sensitivity of Neurones in Aplysia to Temporal Pattern of Arriving Impulses. *Journal of Experimental Biology*, 40(4):643–667, December 1963.
- [119] D. Slepian and H. O. Pollak. Prolate Spheroidal Wave Functions, Fourier Analysis and Uncertainty - I. *Bell System Technical Journal*, 40(1):43–63, January 1961.
- [120] V. K. Somers, M. E. Dyken, and J. L. Skinner. Autonomic and hemodynamic responses and interactions during the Mueller maneuver in humans. *Journal of the autonomic nervous system*, 44(2-3):253–259, 1993.
- [121] C. M. St Croix, M. Satoh, B. J. Morgan, J. B. Skatrud, and J. A. Dempsey. Role of respiratory motor output in within-breath modulation of muscle sympathetic nerve activity in humans. *Circulation research*, 85(5):457–469, September 1999.
- [122] Craig D. Steinback, Aryan Salmanpour, Toni Breskovic, Zeljko Dujic, and J. Kevin Shoemaker. Sympathetic neural activation: an ordered affair. *The Journal of physiology*, 588(Pt 23):4825–4836, December 2010.

- [123] John C. Stevens, Eric P. Lofgren, and Peter J. Dyck. Histometric evaluation of branches of peroneal nerve: Technique for combined biopsy of muscle nerve and cutaneous nerve. *Brain Research*, 52:37–59, March 1973.
- [124] C. O. Tan, J. A. Taylor, A. S. H. Ler, and M. A. Cohen. Detection of multifiber neuronal firings: a mixture separation model applied to sympathetic recordings. *IEEE transactions on bio-medical engineering*, 56(1):147–158, January 2009.
- [125] David J. Thomson. Spectrum estimation and harmonic analysis. *Proceedings of the IEEE*, 70(9):1055–1096, 1982.
- [126] Wilson Truccolo, Uri T. Eden, Matthew R. Fellows, John P. Donoghue, and Emery N. Brown. A Point Process Framework for Relating Neural Spiking Activity to Spiking History, Neural Ensemble, and Extrinsic Covariate Effects. *Journal of Neurophysiology*, 93(2):1074–1089, February 2005.
- [127] R. Tsukahara and T. Mano. The recruitment pattern of single vasoconstrictor neurons in human. *Journal of the autonomic nervous system*, 66(1-2):26–34, September 1997.
- [128] A. B. Vallbo, K. E. Hagbarth, H. E. Torebjörk, and B. G. Wallin. Somatosensory, proprioceptive, and sympathetic activity in human peripheral nerves. *Physiological reviews*, 59(4):919–957, October 1979.
- [129] B. Wallin. Regulation of sympathetic nerve traffic to skeletal muscle in resting humans. *Clinical Autonomic Research*, 16(4):262–269, 2006.
- [130] B. G. Wallin, D. Burke, and S. Gandevia. Coupling between variations in strength and baroreflex latency of sympathetic discharges in human muscle nerves. *The Journal of physiology*, 474(2):331–338, January 1994.
- [131] B. G. Wallin and J. Fagius. Peripheral sympathetic neural activity in conscious humans. *Annual review of physiology*, 50:565–576, 1988.
- [132] Xiaodong Wang, Rong Chen, and Dong Guo. Delayed-pilot sampling for mixture Kalman filter with application in fading channels. *IEEE Transactions on Signal Processing*, 50(2):241–254, February 2002.
- [133] Yiwen Wang, A. R. C. Paiva, and J. C. Principe. A Monte Carlo Sequential Estimation for Point Process Optimum Filtering. In *International Joint Conference on Neural Networks, 2006*, pages 1846–1850. IEEE.
- [134] Yiwen Wang, António R. Paiva, José C. Principe, and Justin C. Sanchez. Sequential Monte Carlo point-process estimation of kinematics from neural spiking

- activity for brain-machine interfaces. *Neural computation*, 21(10):2894–2930, October 2009.
- [135] E. Zaydens, J. A. Taylor, M. A. Cohen, and U. T. Eden. Characterization and Modeling of Muscle Sympathetic Nerve Spiking. *IEEE Transactions on Biomedical Engineering*, 60(10):2914–2924, October 2013.

CURRICULUM VITAE

