

**A Computational Framework for Detecting Rhythmic Spiking Through the Power Spectra  
of Point Process Model Residuals**

by

**Karin Cox**

B.A., Macalester College, 2003

Ph.D., University of Pittsburgh, 2011

Submitted to the Graduate Faculty of the  
Dietrich School of Arts and Sciences in partial fulfillment  
of the requirements for the degree of  
Doctor of Philosophy

University of Pittsburgh

2024

UNIVERSITY OF PITTSBURGH  
DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

**Karin Cox**

It was defended on

December 5, 2023

and approved by

Rebecca Hwa, Professor, Department of Computer Science, The George Washington University

Jacob Biehl, Associate Professor, Department of Computer Science, University of Pittsburgh

Jonathan Rubin, Professor, Department of Mathematics, University of Pittsburgh

Thesis Co-Advisor: Taieb Znati, Professor Emeritus, Department of Computer Science,  
University of Pittsburgh

Thesis Co-Advisor: Robert S. Turner, Professor, Department of Neurobiology, University of  
Pittsburgh

Copyright © by Karin Cox

2024

# **A Computational Framework for Detecting Rhythmic Spiking Through the Power Spectra of Point Process Model Residuals**

Karin Cox, PhD

University of Pittsburgh, 2024

Neurons communicate through rapid, all-or-nothing action potential events (“spikes”). Researchers often predict that oscillatory drives will shape spike trains. For example, computational models of Parkinson’s Disease (PD) predict pathological 12-30 Hz spike rhythms. The detection of spike train oscillations presents an algorithmic challenge. We must devise an automated, scalable means of inferring when a noisy point process arose from a rate function that oscillates at a specific frequency, versus one that oscillates at a different frequency, or does not reliably oscillate.

To identify oscillations, a naïve algorithm might compute the spike train’s power spectral density (PSD) – the distribution of signal power over frequencies – and detect oscillations as significant PSD peaks, contrasted against an assumed flat baseline. Yet non-oscillatory spike trains can exhibit aperiodic features that render this flat baseline inappropriate.

This dissertation investigates whether two common baseline-distorting features can be removed through point process models (PPM), which predict instantaneous spike rates as a function of covariates. I first focus on the “recovery period” (RP): an inevitable, transient post-spike suppression in subsequent spiking. The RP creates global spectral distortion. An established “shuffling” method removes this distortion, but can also reduce the power associated with true spike rhythms. I developed an alternative “residuals” method that accounts for the RP-associated variance in the spike train with a PPM, and generates a corrected PSD from the PPM residuals.

In some spike trains, a second, “burst-firing” feature can create further distortion. To accommodate bursts, I developed a “two-state” residuals method. This method infers the timing of burst- and non-burst states, and separately accounts for these states in the PPM.

I compared the above methods’ ability to enable accurate oscillation detection with flat baseline-assuming tests. Over synthetic data, the residuals method improved upon the shuffling method’s detection accuracy, and the two-state variant offered further improvement when bursts were simulated. Moreover, in empirical data from a parkinsonian monkey, the residuals PSDs yielded increased incidence of the anticipated pathological oscillations.

This work demonstrates that we can use PPMs to remove the distortion that aperiodic features introduce into power spectra, thereby improving the sensitivity and specificity of oscillation detection.

## Table of Contents

<b>1.0 Introduction.....</b>	<b>1</b>
<b>2.0 Background and Related Work.....</b>	<b>10</b>
<b>2.1 Publishing Note.....</b>	<b>10</b>
<b>2.2 Preliminaries .....</b>	<b>10</b>
<b>2.3 Recovery Period.....</b>	<b>11</b>
<b>2.4 Recovery Period Distortion of the Power Spectrum .....</b>	<b>12</b>
<b>2.5 Inter-spike Interval Shuffling Method for Power Spectrum Correction .....</b>	<b>17</b>
<b>2.6 Recovery Period Estimation in Time Domain Methods for Oscillation Analysis ..</b>	<b>19</b>
<b>2.7 Burst Detection: Priorities .....</b>	<b>23</b>
<b>2.8 Burst Detection: Traditional Methods.....</b>	<b>25</b>
<b>2.9 Burst Detection: State Inference Methods .....</b>	<b>27</b>
<b>2.9.1 Recap of Basic HMM Concepts .....</b>	<b>27</b>
<b>2.9.2 Tokdar et al. (2010): HSMM and HMM Approaches to Burst Detection ....</b>	<b>29</b>
<b>3.0 One-State Residuals Method: Materials and Methods.....</b>	<b>33</b>
<b>3.1 Publishing Note.....</b>	<b>33</b>
<b>3.2 Ethics Statement .....</b>	<b>33</b>
<b>3.3 Methods Overview .....</b>	<b>34</b>
<b>3.4 Synthetic Data Methods .....</b>	<b>35</b>
<b>3.4.1 Spike Train Simulation.....</b>	<b>35</b>
<b>3.4.2 Uncorrected PSD.....</b>	<b>36</b>

3.4.3 Shuffling-corrected PSD.....	37
3.4.4 Residuals-corrected PSD .....	38
3.4.4.1 Estimation of Recovery Period Duration .....	38
3.4.4.2 Bounded Last-Spike Point Process Model.....	40
3.4.4.3 Spectral Analysis of Model Residuals .....	43
3.4.4.4 Clarification Regarding “Bounded” TerminologyError! Bookmark not defined.	
3.4.5 Evaluation of the PSD Correction Methods .....	44
3.4.5.1 Synthetic Spike Train Test Sets.....	44
3.4.5.2 Labeling of Hits and False Alarms.....	46
3.4.5.3 Comparison of ROC curves for the Correction Methods .....	47
3.4.5.4 Analysis of Simulation Parameter Effects on Relative Hit and False Alarm Rates.....	51
3.4.5.5 Post-Hoc Analyses of False Alarms in Spike Trains with High Firing Rates.....	51
3.5 Experimental Data Methods.....	53
3.5.1 Subject.....	53
3.5.2 Task .....	53
3.5.3 Surgery .....	54
3.5.4 MPTP Administration .....	54
3.5.5 Data Acquisition and Spike Sorting .....	55
3.5.6 Task-Aligned Analysis Windows .....	56
3.5.7 Unit and Trial Selection Criteria .....	56

3.5.8 Windowed Data: Uncorrected PSD.....	58
3.5.9 Windowed Data: Shuffling-corrected PSD.....	58
3.5.10 Windowed Data: Residuals-corrected PSD .....	60
3.5.11 Reference Synthetic Data for Comparison with Bursting Units.....	61
3.5.12 Population-Level Comparisons of Oscillation Detection Rates .....	61
<b>4.0 One-State Residuals Method: Results.....</b>	<b>63</b>
4.1 Publishing Note.....	63
4.2 Results Overview .....	63
4.3 Synthetic Data Results .....	64
4.3.1 Shuffling Method: Examples of Detected and Non-Detected Oscillations....	64
4.3.2 Residuals Method: Procedure and Example Case Performance.....	66
4.3.2.1 RP Estimation Algorithm.....	67
4.3.2.2 Bounded Last-Spike Point Process Model.....	68
4.3.2.3 Spectral Analysis.....	69
4.3.3 Methods Evaluation .....	70
4.3.3.1 Oscillation Detection Performance: Low-to-Moderate Spike Rates.	70
4.3.3.2 Oscillation Detection Performance: High Firing Rates .....	76
4.3.3.3 Evaluation of Recovery Period Duration Estimation.....	80
4.4 Experimental Data Results .....	82
4.4.1 Motivation and Approach .....	82
4.4.2 Single Unit Examples .....	84
4.4.3 GPI and VLa Population Results.....	91
<b>5.0 Two-State Residuals Method: Materials and Methods .....</b>	<b>93</b>



<b>5.1 Methods Overview</b> .....	<b>93</b>
<b>5.2 Spike Train Simulation: Expansion to Two States</b> .....	<b>94</b>
<b>5.3 Burst detection: Estimation of Initial HMM Settings</b> .....	<b>96</b>
<b>5.3.1 Preliminary Estimation of State-Specific ISI Distributions</b> .....	<b>97</b>
<b>5.3.2 Preliminary Estimation of Initial State Distribution</b> .....	<b>102</b>
<b>5.3.3 Preliminary Estimation of State Transition Matrix</b> .....	<b>102</b>
<b>5.4 Burst detection: Refinement of HMM Parameters</b> .....	<b>105</b>
<b>5.5 Burst detection: HMM-Based Decoding of ISI States</b> .....	<b>106</b>
<b>5.6 Two-State Residuals-corrected PSD</b> .....	<b>108</b>
<b>5.6.1 Estimation of Two-State Recovery Period Durations</b> .....	<b>108</b>
<b>5.6.2 Two-State Bounded Last-Spike Point Process Model</b> .....	<b>109</b>
<b>5.6.3 Spectral Analysis of the Two-State Model Residuals</b> .....	<b>109</b>
<b>5.7 Preliminary Evaluation of the Two-State Residuals Method</b> .....	<b>110</b>
<b>5.7.1 Two-State Synthetic Spike Train Test Sets</b> .....	<b>110</b>
<b>5.7.2 Evaluation of Oscillation Detection Outcomes</b> .....	<b>111</b>
<b>5.7.3 Evaluation of Burst Detection Outcomes</b> .....	<b>112</b>
<b>6.0 Two-State Residuals Method: Results</b> .....	<b>113</b>
<b>6.1 Example #1: Long Non-Burst Recovery Period and Two-Spike Bursts</b> .....	<b>113</b>
<b>6.2 Example #2: Matching Two-State RPs and a Prominent 1/f-like Trend</b> .....	<b>117</b>
<b>6.3 Example #3: Moderate Non-burst RP and Prominent 1/f-like Trend</b> .....	<b>120</b>
<b>6.4 Summaries over Multiple Simulations</b> .....	<b>122</b>
<b>7.0 Conclusions and Future Directions</b> .....	<b>126</b>
<b>7.1 Publishing Note</b> .....	<b>126</b>

<b>7.2 Chapter Overview .....</b>	<b>126</b>
<b>7.3 One-State Residuals Method .....</b>	<b>127</b>
<b>7.3.1 Predicted Findings .....</b>	<b>128</b>
<b>7.3.2 Effect of false alarms with fast spiking and strong modulation .....</b>	<b>130</b>
<b>7.4 Two-State Residuals Method .....</b>	<b>132</b>
<b>7.5 Areas for Future Work .....</b>	<b>134</b>
<b>7.5.1 Further Evaluation of the Burst Detection Step .....</b>	<b>134</b>
<b>7.5.2 Adaptation to Time-Frequency Analysis .....</b>	<b>135</b>
<b>7.6 Concluding Remarks.....</b>	<b>137</b>
<b>Appendix A Supporting Information, Figures and Tables for the One-State Residuals</b>	
<b>Method .....</b>	<b>138</b>
<b>Appendix A.1 Shuffling Correction by Division: Rationale and Assumption Evaluation</b> <b>.....</b>	<b>138</b>
<b>Appendix A.2 Method Evaluation on an Alternative Simulation Framework.....</b>	<b>142</b>
<b>Appendix A.2.1 Simulation Methods.....</b>	<b>142</b>
<b>Appendix A.2.2 Simulation Results .....</b>	<b>148</b>
<b>Appendix A.3 Supporting Figures .....</b>	<b>154</b>
<b>Appendix A.4 Supporting Tables.....</b>	<b>168</b>
<b>Appendix B Supporting Figures and Tables for the Two-State Residuals Method .....</b>	<b>175</b>
<b>Appendix B.1 Supporting Figures.....</b>	<b>175</b>
<b>Bibliography .....</b>	<b>182</b>

## List of Tables

<b>Table 1. Parameters for Simulation of Two Additional Sets of Bursty Units. ....</b>	<b>111</b>
<b>Table 2. Rates of hit-containing and false alarm (FA) containing PSDs for the “bursty” units. .....</b>	<b>124</b>
<b>Table 3. Performance of “bursty” versus “nonbursty” classification for the two-state residuals method. ....</b>	<b>124</b>
<b>Table 4. For the three parameter sets, the accuracy, hit (second row), and false alarm (third row) rates for the labeling of ISIs within individual spike trains.....</b>	<b>125</b>
<b>Table A1. Parametric Results for the Methods' Hit Rate Differences: Low-to-Moderate Firing Rates. ....</b>	<b>169</b>
<b>Table A2. Parametric Results for the Methods' False Alarm Rate Differences: Low-to- Moderate Firing Rates.....</b>	<b>170</b>
<b>Table A3. Parametric Results for the Methods' False Alarm Rate Differences: Low-to- Moderate Firing Rates.....</b>	<b>171</b>
<b>Table A4. Parametric Results for the False Alarm Differences: Manipulation of Base Firing Probability. ....</b>	<b>172</b>
<b>Table A5. Parametric Results for the Hit Rate Differences: High Firing Rates.....</b>	<b>173</b>
<b>Table A6. Parametric Results for the False Alarm Differences: High Firing Rates. ....</b>	<b>174</b>

## List of Figures

<b>Figure 1. Examples of Recovery Period Distortion of Non-Oscillating Power Spectra. ....</b>	<b>14</b>
<b>Figure 2. Distortion of Oscillating Power Spectra, and Illustration of the Shuffling Method. .....</b>	<b>16</b>
<b>Figure 3. Distortion correction by the proposed residuals method.....</b>	<b>42</b>
<b>Figure 4. Performance of the shuffling and residuals methods over a diverse synthetic dataset. ....</b>	<b>49</b>
<b>Figure 5. Residuals (res) - shuffling (shuf) difference in hit rates (<math>D_{HR}</math>) over the varied parameters of the primary synthetic dataset.....</b>	<b>73</b>
<b>Figure 6. Residuals - shuffling difference in false alarm rates (<math>D_{FA}</math>) over the varied parameters of the primary synthetic dataset.....</b>	<b>75</b>
<b>Figure 7. Performance of the shuffling and residuals methods over a synthetic dataset of high firing rate (FR) spike trains. ....</b>	<b>77</b>
<b>Figure 8. Residuals - shuffling difference in hit rates over the varied parameters of the dataset of high firing rate spike trains. ....</b>	<b>79</b>
<b>Figure 9. Residuals - shuffling difference in false alarms over the varied parameters of the dataset of high firing rate spike trains. ....</b>	<b>81</b>
<b>Figure 10. Comparison of the shuffling and residuals output for two highlighted units acquired from a parkinsonian non-human primate (NHP).....</b>	<b>85</b>
<b>Figure 11. Comparison of the shuffling and residuals output for two units with putative beta oscillations and non-oscillatory features. ....</b>	<b>87</b>

<b>Figure 12. Points of significant spectral power identified by the shuffling and residuals methods across the population of analyzed spike trains acquired from the parkinsonian non-human primate (NHP). .....</b>	<b>89</b>
<b>Figure 13. Demonstration of the procedure for separation of an ISI distribution into burst- and non-burst components.....</b>	<b>100</b>
<b>Figure 14. True and estimated ISI probabilities for the synthetic approximation of the empirical unit illustrated in Figure 11A. ....</b>	<b>115</b>
<b>Figure 15. Comparison of the shuffling output, and the one-state (1S), and two-state (2S) residuals output for the unit simulated to match the empirical unit from Figure 11A. ....</b>	<b>116</b>
<b>Figure 16. True and estimated ISI probabilities for the synthetic approximation of the empirical unit illustrated in Figure 11B.....</b>	<b>118</b>
<b>Figure 17. Comparison of the shuffling, one-state (1S) residuals output and two state (2S) residuals output for the unit simulated to match the empirical unit from Figure 11B. ....</b>	<b>119</b>
<b>Figure 18. True and estimated ISI probabilities for a synthetic approximation for a synthetic unit with opposite-direction RP and burst effects on the shortest lags in the hazard function. ....</b>	<b>121</b>
<b>Figure 19. Comparison of the shuffling, one-state (1S) residuals output and two state (2S) residuals output for the unit simulated with a moderate-length non-burst RP duration.....</b>	<b>123</b>
<b>Figure A1. Simulation-based assessment of the assumptions of the proposed proof from Rivlin-Etzion et al. [11].....</b>	<b>140</b>

<b>Figure A2. Output of the three PSD methods (uncorrected, shuffling, and residuals) on the first example spike train generated under the alternative simulation framework.</b>	<b>149</b>
<b>Figure A3. Output of the three PSD methods on the second example spike train generated under the alternative simulation framework. ....</b>	<b>150</b>
<b>Figure A4. Output of the three PSD methods on the thrd example spike train generated under the alternative simulation framework. ....</b>	<b>153</b>
<b>Figure A5. Estimation of recovery period duration: No oscillation example.....</b>	<b>154</b>
<b>Figure A6. Residuals - shuffling difference in hit rates over the varied parameters of the dataset formed with direct base firing rate manipulation. ....</b>	<b>155</b>
<b>Figure A7. Residuals – shuffling difference in false alarm rates over the varied parameters of the dataset formed with direct base firing rate manipulation.....</b>	<b>156</b>
<b>Figure A8. Residuals and shuffling hit and false alarm rates for the dataset of high firing rate spike trains, with the oscillation frequency (<math>f_{osc}</math>) fixed at 7 Hz. ....</b>	<b>157</b>
<b>Figure A9. Residuals and shuffling hit and false alarm rates for the dataset of high firing rate spike trains, with the oscillation frequency (<math>f_{osc}</math>) fixed at 32 Hz. ....</b>	<b>158</b>
<b>Figure A10. Example power spectra and spike train segments for high firing rate cases that produced pronounced false alarms following residuals or shuffling correction.....</b>	<b>159</b>
<b>Figure A11. Example residuals-corrected power spectra, and corresponding spike train samples, for two levels of oscillation frequency (<math>f_{osc}</math>). ....</b>	<b>161</b>
<b>Figure A12. Performance of the shuffling and residuals methods over a synthetic dataset generated using a short, absolute recovery period. ....</b>	<b>163</b>
<b>Figure A13. Performance of the shuffling and residuals methods over a synthetic dataset generated using a shortened relative recovery period.....</b>	<b>164</b>

<b>Figure A14. Performance of the shuffling and residuals methods over a synthetic dataset generated using a lengthened relative recovery period. ....</b>	<b>165</b>
<b>Figure A15. Residuals and shuffling hit and false alarm rates for the primary synthetic dataset, with the oscillation frequency (<math>f_{osc}</math>) fixed at 12 Hz. ....</b>	<b>166</b>
<b>Figure A16. Inter-spike interval distributions and RP duration estimation for the two VLa units with suspected oscillatory and non-oscillatory features. ....</b>	<b>167</b>
<b>Figure B1. Illustration of the procedure used to obtain an initial estimate of the state transition matrix, <math>A</math>, to be applied to the initialization of the Hidden Markov Model (HMM). ....</b>	<b>176</b>
<b>Figure B2. Demonstration of the ISI separation procedure, for the simulations informed by the empirical unit in Figure 11B.....</b>	<b>178</b>
<b>Figure B3. Comparison of the shuffling, one-state (1S) residuals output and two state (2S) residuals output for the non-bursty control unit, matched to the simulation designed to approximate the empirical unit from Figure 11B.....</b>	<b>180</b>
<b>Figure B4. Demonstration of the ISI separation procedure, for parameter settings expected to generate a rise in the non-burst hazard function over the shortest ISIs (for the “bursty” unit only).....</b>	<b>181</b>

## 1.0 Introduction

Considered at the most abstract level, this thesis addresses a challenge in the recognition of patterns in time series. These patterns likely appear across diverse domains, although I focus on the context of neuroscientific data. The patterns are considered interesting in their own right, and can inform practical applications.

In the neural domain, the time series of interest – the “spike train” -- reports the timestamps of a single neuron’s action potentials. Action potentials (the “spikes”) consist of an all-or-nothing, brief ( $\approx 1$  ms), and largely stereotyped fluctuation in a neuron’s membrane potential [1]. Their occurrence initiates the communication between connected units. Specifically, spikes trigger the release of neurotransmitters into the synaptic cleft, which in turn may promote or inhibit spiking in downstream neurons. Neural models typically assume that the timing of spike events (as opposed to variability in waveform shape) encodes the information transmitted between individual units, thereby justifying their representation as point processes of spike arrival times.

In spike trains, the patterns that we seek are oscillations. For this thesis, I will adopt a working definition of an oscillation as a narrowband, and approximately sinusoidal modulation of a time series. In the spike train setting, the modulation applies to the instantaneous rate function that gave rise to the observed spikes [2].

Why is it important to recognize oscillations in neuronal spiking? Since this thesis considers data from research on Parkinson’s Disease (PD), I will offer two answers specific to this field. First, accurate spike oscillation estimates can inform ongoing development of computational models of PD pathophysiology [2]. Existing models already posit a role for oscillations in various



frequency bands (e.g., the 8-12 Hz “alpha” band, or the 12-30 Hz “beta” band) in driving PD motor symptoms [3]. However, these modeling efforts often must rely on oscillation estimates obtained from coarser signals (e.g., local field potentials, or “LFPs”, which reflect a complex mixture of signal contributions from populations of neurons [4]). Such models could arguably be improved by increased access to spike oscillation estimates, given the fine-grained information these would provide regarding the core unit of neuronal communication. Second, accurately-recognized spike oscillations could serve an important role in adaptive deep brain stimulation (DBS) therapy for PD. In contrast to conventional DBS, which uses a constant pulse sequence, novel adaptive DBS protocols exert real-time, closed-loop control of the stimulation pattern, as determined by algorithms that analyze the features of incoming neural recordings. Many adaptive protocols use oscillations as a major source of control feedback [5]. Again, the estimates are typically derived from coarse population-level measures, raising the possibility that the more precise spike oscillation information could improve this investigational technology.

The primary objective of this thesis is to develop accurate computational methods for identifying oscillations in spike trains. To be more specific, I will cast this objective as a set of algorithmic requirements. The input that an oscillation-detection algorithm must process is a sequence of timestamps of spike arrivals. Frequently, researchers convert these time stamps into a binary, 1 kHz vector, with the “1s” marking time bins of spike occurrence [6]. This input represents a series of stochastic, biological events, the occurrence of which may be shaped by both a latent rate function and noise. An algorithm must translate the input into an output that provides us with at least three types of information: (1) whether oscillations are likely to have shaped the occurrence rate of the observed spikes, (2) the frequencies at which any identified oscillations occurred, and (3) a quantity that communicates the risk that any labeled oscillations may have

actually been generated by noise, under some alternative, baseline model in which the oscillation is absent (e.g., a  $p$  value). As with the design of any algorithm, we also want to maintain reasonable limits on runtime and memory complexity, although the focus of this thesis is primarily on the accuracy question.

Why is it challenging to meet this objective? A central challenge concerns the choice of the baseline, null hypothesis model. The reason this choice is difficult can be illustrated by considering a basic oscillation detection strategy based on standard signal processing methods. Since we are seeking approximately sinusoidal oscillations, a logical option is to compute the power spectral density (PSD) function for the binary time series. The PSD reflects the squared magnitude of the Discrete Fourier Transform (DFT) of the input (see Chapters 2-3 for details), and therefore, indicates the relative weights assigned to a basis set of sinusoids (spaced at harmonic frequencies) that have been fit to the signal [6]. In the case of unstructured, Poisson spiking (implying a constant rate function), power should be uniformly distributed over frequencies, resulting in a flat spectrum. We might be inclined, then, to choose the flat spectrum as our baseline model, representing the non-oscillatory case. If a spike train is indeed shaped by oscillatory trends, these will emerge as peaks against this otherwise flat baseline. Therefore, one might further propose an oscillation detection algorithm that seeks peaks that rise above a flat baseline, beyond some estimate of the noise in the spectral estimate.

The problem with this simple algorithm is that it neglects the contribution of aperiodic components to the spike rate function, and consequently, the PSD. The DFT fits the aperiodic trends with combinations of sinusoids that are unevenly distributed over frequencies, thereby producing points of elevated spectral power that should not be attributed to an oscillation.

How do we address this problem? The strategy adopted by this thesis, as in prior work [7], is to attempt remove the contributions that anticipated aperiodic components make to the PSD, to reduce the possible non-oscillatory explanations for any departures from the flat baseline model. In neuroscience – as in other domains – we can draw upon pre-existing knowledge of the modeled time series to identify the highest priority aperiodic features to address. In this thesis, I address two such aperiodic features.

The first feature – the recovery period – creates the primary analysis challenge that I address in this work. Here (and consistent with past usage; e.g, [8]) I use the term “recovery period” (RP) to refer to the full interval of transiently suppressed firing that reliably follows a spike. This suppression may be absolute ( $p(\text{spike}) = 0$ ) or relative (gradually recovering  $p(\text{spike})$ ). An RP includes, at minimum, an initial refractory period (a typically 5-10 ms [1] interval during which well-known ion channel dynamics prevent spiking), which in some units may be followed by continued suppression due to more variable mechanisms (e.g., circuit-level neuronal interactions). The inevitability of at least a short RP motivates the high priority placed on addressing this aperiodic feature.

The recovery period introduces an inevitable, negative short-lag autocorrelation in the spike rate function, resulting in a well-characterized, broadband distortion of a non-oscillatory PSD, away from the flat ideal [7, 9, 10]. In Chapter 2, I discuss this distortion pattern in detail, and the oscillation detection errors that commonly result when researchers inaccurately assume a null model of Poisson spiking, without addressing the influence of the RP.

Some existing methods do attempt to control for the RP when assessing oscillations, but present limitations that would be useful to address. Of these existing methods, I will primarily focus on the inter-spike interval (ISI) “shuffling” method [7]. The shuffling method adopts the

corrective strategy described above: The authors attempt to model and remove the undesirable contributions to the PSD, to increase confidence that peaks above a flat baseline do reflect spike rate oscillations. I will argue that the shuffling method can perform effectively in many scenarios, but in some cases (e.g., slowly-spiking units) removes temporal structure that is critical for the detection of true spike rhythms. Ideally, one could devise an alternative approach that spares this critical information. Another category of methods, which replace the PSD with time domain models ([11-13]), suggest some strategies to achieving such sparing, but I will argue that shortcomings hinder their practical adoption. Two such shortcomings include the lack of the conventional PSD output (which is arguably better suited for capturing the assumed sinusoidal oscillations, and is also of historical importance in neuroscience), and these methods' dependence on manually-set free parameters, for which the appropriate settings may not be obvious. Overall, we are left with a need for an RP removal method that spares important oscillatory information, while minimizing vulnerability to arbitrary configuration decisions, and providing a standard spectral representation.

The second aperiodic feature – burst-firing -- may drive further departure from the flat baseline model, at least in some units. Although neuroscientists use the term “burst” somewhat liberally and flexibly [14], the label typically refers to abrupt increases in spiking beyond a background rate [15], consistent with a switch to a distinct phase of spike generation dynamics [16]. As I will further discuss and demonstrate in Chapters 5-6, this switch can further complicate oscillation detection in two ways: First, burst and non-burst states may feature differing RP dynamics, thereby implying state-dependent RP distortion of the spike train's frequency content. Second, the alternating firing rates contribute to a  $1/f$ -like trend (i.e., a gradual decay in power from the low to high frequencies) in the power spectrum [17].

As a strategy to remove these spectral components, one might attempt to detect burst- and non-burst states in spike trains, and estimate and control for the states' distinct spike rate and RP properties. Several burst detection algorithms exist (see [14] for reviews), and two of these (both reported in [18]) approach burst detection as a state inference problem, returning both a time series of state labels and expected spiking properties associated with those states. However, as I discuss further in Chapter 2, the practical deployment of this state inference approach would benefit from improvements upon the currently available options. For example, and similar to the time-domain RP-correcting methods mentioned above, the existing burst-state inference methods require the user to set a number of free parameters, which in practice may often be set randomly, or left fixed to a set of hard-coded defaults. I will argue that reliance on such parameter settings may result in poor performance in some cases, thereby raising the need for a burst inference method that reduces the configuration burden.

In light of the above, I supplemented the primary focus of this dissertation (removing the distortion that the RP introduces to the PSD), with an exploratory assessment of a novel burst state inference algorithm, and the use of the inferred states to control for the bursts' impact on the power spectrum. A set of three shared ideas guided both the primary and exploratory efforts. I began with the idea – which drew upon precedents in the above-mentioned time-domain methods [12, 13] -- that point process models (PPMs, which predict instantaneous spike rates as a function of covariates) could account for the RP- and burst-related variance. To this pre-existing idea, I added two new hypotheses. First, I considered that we could effectively detect oscillations by applying spectral analysis to the PPM residuals, and testing for significant peaks against the null hypothesis of a flat spectrum. Second, I considered that the PPMs could be configured in an automated, data-

driven, and neuron-specific manner, primarily on the basis of measures derived from the distribution of inter-spike intervals (ISIs).

To assess these hypotheses, I modeled the RP effects by applying a Poisson Generalized Linear Model (GLM) to each individual unit's spike train. As described in detail in Chapter 3, each GLM modeled time-varying spiking intensity as a function of the time transpired since the most recent spike's occurrence, provided that this spike occurred within a limited window of spike history. The width of the history window – which, ideally, would match the typical duration of a unit's recovery period – was determined by an algorithm that estimated RP duration on the basis of the distribution of each unit's ISIs. Each PPM's residuals were submitted to spectral analysis, and putative oscillations were identified as power spectrum peaks that exceeded a flat null hypothesis model.

This initial method (termed the “residuals method”) was evaluated on both a diverse set of synthetic spike trains (for which ground truth oscillations were known) and empirical data that had been acquired from a monkey model of Parkinson's Disease (in which pathological oscillations were anticipated; see Chapter 4). Results were compared against those derived from the standard ISI shuffling procedure. Across the majority of the synthetic spike trains, the residuals method met or improved upon the shuffling method's accuracy in oscillation detection. In the empirical data (with acknowledgment of the lack of ground truth knowledge in this case), the trends were consistent with enhanced residuals' method sensitivity to the expected pathological oscillations, especially in units with low firing rates (FRs). At the same time, individual instances in the empirical data did point to the complications that remain when the PPM does not account for phenomena such as bursting.

To explore the feasibility of expanding this general strategy to accommodate bursting, I implemented a new PPM-based method that was informed by estimates of spike trains' burst and non-burst states. This method is described in detail in Chapter 5. Briefly, I devised an algorithm to estimate, from each spike train's ISI histogram, rough approximations of the ISI distribution statistics required for a Hidden Markov Model (HMM), which would in turn infer the likelihood that each ISI emerged from either a burst or non-burst state. These rough approximations were further refined through the standard algorithm for unsupervised training of an HMM (the Baum-Welch algorithm), and the final HMM parameters were used to output the most likely sequence of states visited by the spike train. These state labels – along with state-specific RP duration estimates, derived from state-specific ISI distribution estimates – were used to build two-state PPMs that separately modeled baseline spike rate and RP effects for the inferred burst and non-burst intervals. As with the original PPM approach described above, the model residuals were submitted to spectral analysis, and significant peaks were sought against the null hypothesis of a flat spectrum.

As this expansion of the residuals method is an exploratory development, a detailed evaluation has yet to be completed. However, in Chapter 6 I demonstrate an initial proof-of-concept of the approach, by showing how the switch from the one-state to the two-state method improves the removal of the RP distortion and  $1/f$ -like trends from the power spectrum, within sample synthetic units that were simulated to generate both bursting and oscillations.

In summary, this thesis presents the following contributions:

- (1) Development of a computational framework – the “residuals method” -- for removing the influence of the recovery period on a spike train's power spectrum, through application of point process models tailored to the individual spike train instances.

- (2) Expansion of the original framework to detect when the spike train is best described as alternating between “burst” and “non-burst” states, and to construct PPMs that accommodate the differing recovery periods that may occur during these states.
- (3) Development of a strategy for evaluating the original and expanded frameworks on synthetic spike trains.
- (4) Demonstration, through implementation of (3), of (a) the superior performance of the original residuals method, as compared to an established alternative, and (b) the further improvements produced by the expanded residuals method.
- (5) Application of the original residuals method to experimentally acquired data from real neurons, allowing demonstration of the method’s potential for enhancing sensitivity to the oscillations expected in this setting.

In the remainder of this thesis, I begin with a review of the related literature in Chapter 2. Subsequently, in Chapter 3, I describe the design of the original, one-state residuals method. In Chapter 4, I describe this method’s extensive evaluation, as compared to the performance of the shuffling method. In Chapter 5, I describe the design of the exploratory, two-state method, and in Chapter 6 I demonstrate its performance, relative to the one-state method, on sample synthetic units. Chapter 7 summarizes the major conclusions and outstanding questions to be pursued in future research.



## **2.0 Background and Related Work**

### **2.1 Publishing Note**

Some text and figures in this Chapter have been adapted from a manuscript that has been posted to a preprint server [19] and submitted to a peer-reviewed journal.

### **2.2 Preliminaries**

This section begins with three brief explanations of conventions and terms which may not necessarily be familiar to readers from fields outside of neuroscience.

First, I will clarify a few details regarding single-unit spike train time series. Both spike trains and local field potentials (LFPs) are derived from extracellular potential signals, which have been recorded from individual channels of implanted electrodes. As mentioned in the Introduction, LFPs reflect aggregate neural activity, and result from the application of a low pass filter to the voltage trace. In contrast, an individual spike train indicates the timestamps of action potentials that have been detected in this same voltage signal, and attributed to a single source neuron. The spike trains are identified through a combination of manual and/or automated “spike sorting” procedures (see [20] for a recent review). Each spike train is typically represented as a 1 kHz binary time series, where a “1” marks a millisecond-width time bin in which a spike was inferred to have been generated by the single unit.

Second, the spectral analysis of spike trains typically entails the direct submission of this binary time series (post mean-removal) to a standard procedure (e.g., Welch's method or the multitaper method) that combines the Discrete Fourier Transform with appropriate windowing and smoothing steps (see [6] for a review, and Chapter 3 for details regarding the specific method used in the current research). As briefly mentioned in the Introduction, neuroscientists conventionally divide the power spectrum into a series of physiologically-relevant frequency bands [21]: delta (0.5-4 Hz), theta (4-8 Hz), alpha (8-12 Hz), beta (12-30 Hz), and gamma (> 30 Hz).

Third, a few additional statements regarding the neuronal refractory period will provide useful background for the Recovery Period subsection below. As briefly mentioned in the Introduction, the neuronal refractory period is a mandatory event, which entails the immediate post-spike suppression of new spike generation, through changes in the status of ion channels in the neuron's membrane. The reader is referred to general texts for full details of this phenomenon (e.g., [1, 16, 22]). Briefly, note that the refractory period is traditionally divided into absolute and relative stages. The absolute refractory period, when empirically estimated, has been shown to last a millisecond or slightly less [12], and refers to an interval when Na<sup>+</sup> and K<sup>+</sup> channel states render a repeat spike impossible. Subsequently, as the ion channels gradually return to their pre-spike status, the neuron proceeds through the relative refractory period, when spiking is possible, but requires increased external drive to occur.

### **2.3 Recovery Period**

The majority of the analysis procedures that I review describe the estimation of a refractory period, rather than recovery period effect [11-13]. Therefore, it is necessary to clarify my choice of the latter,

mechanism-neutral term. In data from our lab (including the empirically-sampled units we present here), we have observed several examples of units that exhibit a dead time (i.e., sharp post-spike drops in the spike hazard function) that far exceeds the typical duration of the channel dynamics responsible for refractoriness (as mentioned in the Introduction, approximately 5-10 ms for the absolute and relative phases combined [1]). Such long post-spike pauses might be generated by either circuit-level or intrinsic mechanisms (such as those underlying pacemaker patterns [23]). Since none of the methods that I discuss in this thesis can claim to selectively correct for the biological refractory period *per se*, I adopt the generic recovery period (RP) terminology.

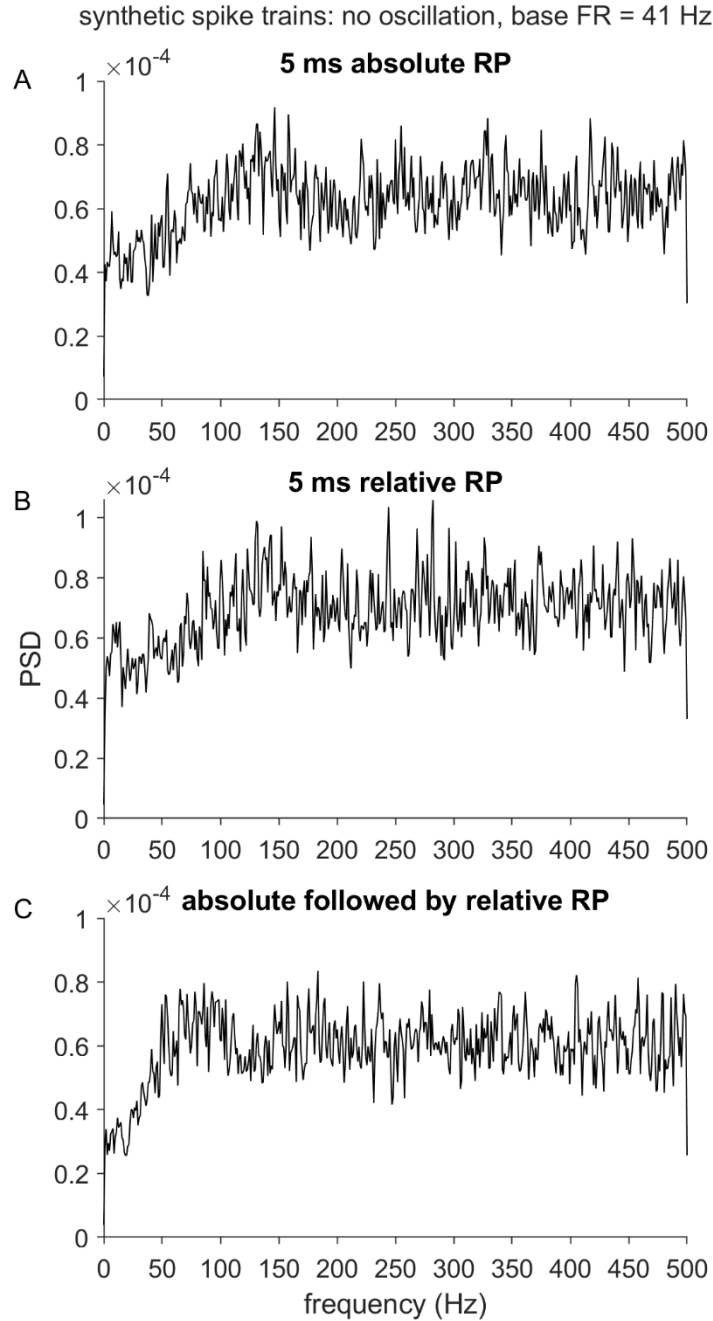
## 2.4 Recovery Period Distortion of the Power Spectrum

In this section I will describe and illustrate the recovery period's distortion of the power spectrum, first in the context of otherwise Poisson spiking, and subsequently in the context of oscillatory spiking. As background, note that concern regarding the challenges created by this general form of distortion is shared beyond neuroscience. For example, in physics and astronomy, the same spectral distortion pattern arises from measurement dead time, referring to the interval after detection of an event (e.g., a photon) when the processing requirements of a measurement device prevent detection of a second event [24]. However, the fact that this dead time arises from controllable and human-designed devices implies that it can be addressed by correction strategies that are not usable by researchers studying neurons (e.g., because the correction strategy assumes the dead time duration is fixed and known [25]).

Regardless of the origin, the distortion patterns created by post-event pauses share a set of predictable features. Returning to spike train terminology, we can begin with the simplest illustration of recovery period distortion: a hypothetical spike train that arises from an otherwise Poisson process, apart from post-spike interruptions by a brief absolute recovery period. In this case, the resulting PSD departs from the flat spectrum anticipated for an unmodified Poisson process, and is instead characterized by a trough over a range of low frequencies, followed by elevated power (often accompanied by visible ringing) over the high frequencies ([9, 10]; see Figure 1A for a synthetic example). The precise form of the distortion can differ modestly for relative recovery periods (see Figure 1 for synthetic examples using absolute and relative RPs, and an absolute-relative RP sequence), and is exacerbated by high firing rates (FR) and long recovery period durations [7].

When a spike train features an oscillation, the corresponding peak in the PSD is superimposed on the recovery period-associated distortion pattern (see Figure 2A for synthetic examples; note the remainder of Figure 2 will be discussed in Chapter 4). The distortion complicates standard statistical procedures that aim to distinguish such oscillatory PSD peaks from baseline by using a null hypothesis model that is estimated from the spectrum itself. In particular, detection problems arise when the baseline does not account for the RP contribution.

An especially detailed treatment of RP-associated distortion [7] -- which also introduced the established shuffling method -- illustrated the oscillation detection challenges with a straightforward statistical test that assumes a Poisson process (in the time domain) and corresponding white noise (in the frequency domain) as the baseline. These authors observed that, over the highest frequencies (beyond the frequency bands typically prioritized by neuroscientists),



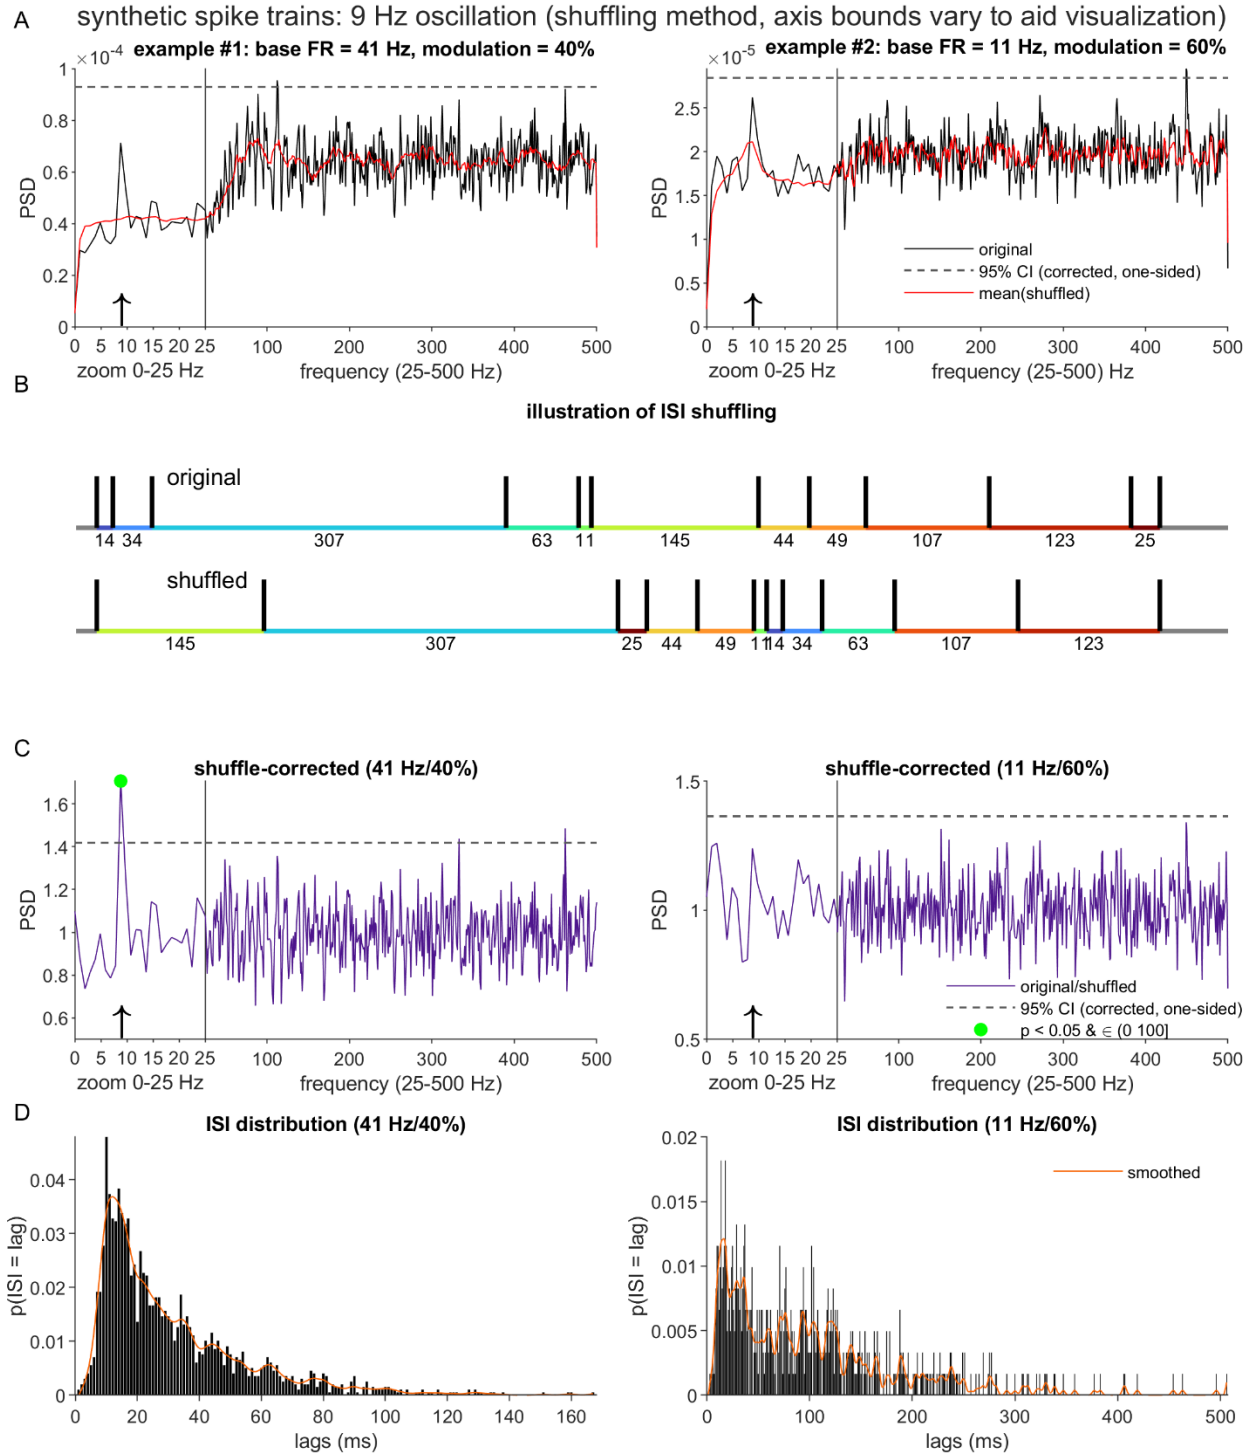
**Figure 1. Examples of Recovery Period Distortion of Non-Oscillating Power Spectra.** The panels depict the power spectral density (PSD) of three spike trains simulated under the assumption of Poisson spiking interrupted by three different forms of recovery period (RP, the post-spike reduction in spike probability, which can extend beyond the refractory period). Relative to the flat spectrum of a Poisson process, all three of the RP-distorted PSDs exhibit a characteristic trough over the lowest frequencies, and relative elevation and ringing over the higher frequencies. All PSDs were generated according to the simulation framework from [11] (see Chapter 3; baseline firing rate (FR) = 41 Hz; simulation duration  $T = 60 \times 1024$  ms; RP duration  $n_r$  and steepness  $k$  parameters vary by panel). (A) PSD of a spike train simulated with an absolute RP ( $n_r = 5$  ms;  $k = 0$ ). (B) PSD of a spike train

simulated with a relative RP ( $n_r = 5$  ms;  $k = 0.5$ ). (C) PSD of a spike train for which each spike triggered an absolute RP, immediately followed by a relative RP. The absolute and relative components adopted the  $n_r$  and  $k$  parameters described for panels (A) and (B), respectively.

the PSD functions of neuronal spike trains tended to approach those of matched-FR Poisson processes, thereby motivating the construction of confidence intervals (CI) based on the mean and variance over these high-frequency points.<sup>1</sup> Points in the PSD function that cross the CI bound (see Figure 2A, dashed lines, for CI examples) are labeled as significant, with the possible added requirement that the points also fall within an *a priori* search range (set at (0, 100] Hz in the figures). Because this statistical test does not account for the RP distortion, it is prone to miss oscillatory peaks that fall within the low frequency trough (see the two examples in Figure 2A), and incorrectly label as significant points at higher frequencies, beyond the trough, which exhibit elevated power as a consequence of the distortion effect alone [7]. Although the examples in Figure 2A do not show false positives within the (0, 100] Hz search range, the PSDs in both panels do include peaks that cross the CI threshold at frequencies  $> 100$  Hz. Note that this general pattern of false negatives and positives also appears when the CIs for the Poisson null hypothesis are estimated using an alternative analytical procedure (see [7] and [26] for details).

---

<sup>1</sup> In the present work, the CI-defining, “control” frequency range will be defined as 250-500 Hz. The 250 Hz lower bound falls well above 100 Hz (the frequency at which [11] observed that PSDs tended to flatten, and a typical upper bound given for the gamma band), and the 500 Hz upper bound is the Nyquist frequency for a 1 kHz time series.



**Figure 2. Distortion of Oscillating Power Spectra, and Illustration of the Shuffling Method.** The panels depict the recovery period distortion of the PSDs for synthetic spike trains with oscillations, and the limitations of inter-spike interval-based correction. (A) PSDs (black lines) for two synthetic spike trains, generated using matching simulation durations ( $T = 120 \times 1024$  ms) and recovery period parameters ( $n_r = 9$  ms;  $k = 0.7$ ), and added oscillation terms of matching oscillation frequency ( $f_{osc} = 9$  Hz, highlighted by arrows on the frequency axis). The spike trains differed with respect to the oscillation modulation strength

( $m$ ) and the baseline firing rate (base FR; see column titles). Dashed lines mark the corrected 95% one-sided confidence interval (CI) constructed under an approximation of a null model of Poisson spiking (see Chapter 3). The testing procedure for statistically significant power (PSD points that both cross the CI bound and fall within an a priori search range of (0, 100] Hz) found no significant points within either spectrum. Red lines plot the spectral component that the ISI shuffling algorithm [7] estimates to be attributable to the structure present in the inter-spike intervals (ISIs). (B) Illustration of one iteration of ISI shuffling on a truncated 1000 ms dataset (actual implementation applied global shuffling to the full spike train duration). The average PSD of several shuffled spike trains forms the estimated ISI-attributable spectrum. (C) The shuffling-corrected PSDs, equal to the ratio of the original and shuffling-estimated PSDs. The green dot (*left*) indicates a PSD point that both crosses the CI bound and falls within an a priori search range of (0, 100] Hz. (D) ISI probability distributions for the two example spike trains. Smoothing spline fits (orange lines) highlight the more visible shaping of the ISI distribution by the 9 Hz oscillation in the lower-FR case. Note that the concurrent shaping of the ISIs by a high frequency rhythm, which is readily visible in both the higher- and lower-FR examples, is attributable to the 9 ms relative recovery period.

## 2.5 Inter-spike Interval Shuffling Method for Power Spectrum Correction

To account for the RP distortion, the correction method proposed by [7] builds upon the null hypothesis that the spike train emerges from a renewal process with a matched inter-spike interval (ISI) distribution. Recall that a renewal process is a sequence of events (here, spikes) for which the waiting times (the ISIs) are independently sampled from an identical distribution [27]. Since we define the recovery period as an influence solely on the waiting time between two successive events, and not on any higher order structure (e.g., correlations between successive waiting times), the renewal null hypothesis does capture the temporal structure attributable to the RP.

The ISI shuffling method described by [7] estimates the renewal equivalent spectrum for a spike train. According to the simplest, global version of this scheme, several control spike trains are generated by randomly permuting the ISIs of the original. The PSDs computed from these



control spike trains are then averaged to estimate the renewal equivalent spectrum. An analytical alternative to this Monte Carlo procedure does also exist, which leverages an established method to compute the renewal process spectrum from an ISI distribution [28, 29]. Although this analytical approach is fast and non-stochastic, it does require the use of rectangular spectral analysis windows. For the present work, we focus on the shuffling approach, given its flexibility to tapered windows (e.g., Hamming windows, or the Slepian sequences used by the multitaper method [30]).

Regardless of whether shuffling or the analytical method is used, the resulting control spectrum is then divided out from the original spectrum to obtain the corrected PSD (see Appendix A.1 for comments regarding the use of division as opposed to subtraction). Given a corrected spectrum, it is customary to proceed with the aforementioned procedure of constructing CIs based on the high frequencies, under the assumption that the null hypothesis of a flat spectrum is now appropriate.

Shuffling correction continues to see common use in single-unit studies, especially in research on the motor control circuitry (e.g., [31-35]), but also in other domains [36, 37]. The method is a well-justified option when researchers are willing to forgo the temporal information present in the ISI distribution. However, as noted by [7], the ISIs can carry information about the rhythmic occurrence of individual spikes, which the shuffling method will remove. A simple, rigid example of this occurs in the case of precisely timed pacemaking activity – that is, single spikes separated by a relatively consistent period  $P$ , plus or minus a modest amount of jitter (i.e., an approximation

of a Dirac comb [38])<sup>2</sup>. In such cases, the pacemaking interval is indistinguishable from a recovery period, implying that such patterns should be detected through means other than RP-corrected spectra (e.g., by seeking instances of extremely low ISI variance). A more subtle example occurs when rhythmic spiking appears to be driven by an approximately sinusoidal function, but the mean FR is so low that spikes often occur only once per oscillation cycle, and may at times skip cycles altogether. To detect and statistically validate such cases, spectral analysis may prove more helpful, but how to effectively address RP distortion while preserving the sparse information regarding the underlying rhythmicity remains unresolved. Framed in computational terms, the key challenge is to identify an appropriate model that more precisely captures the undesirable RP-associated variance.

## **2.6 Recovery Period Estimation in Time Domain Methods for Oscillation Analysis**

As indicated in the Introduction, this thesis will use point process modeling, in the time domain, to address RP distortion, and then perform spectral analysis on the PPM residuals. This approach was informed by preceding methods that estimated oscillatory trends – while also accounting for RP effects -- using PPMs alone. These methods are briefly summarized in this

---

<sup>2</sup> To clarify why shuffling removes evidence of the pacemaking rhythm: If the ISIs are nearly uniform, then the shuffled spike trains are nearly identical to the original spike train. Therefore, the mean of these shuffled spike trains' spectra will be nearly identical to the original spectrum, such that the division-based correction will result in a nearly flat corrected PSD.

section. Relative to ISI shuffling, all of these methods are distinguished by their relative sparing of the temporal structure expressed in long-duration ISIs when controlling for the variance attributable to the recovery period.

In general, spike train PPMs seek to explain time-varying spike probability as a function of a set of covariates, through a variety of modeling strategies (see [6, 8], and a more detailed, technical description in Chapter 3). The PPMs that inform the present work fall into two subcategories. The first subcategory includes methods that, like this thesis, use Poisson regression (see Chapter 3 for details) to predict spike likelihood as a function of a set of covariates informed by the unit's recent spiking history (see [12, 13] for representative examples). These pre-existing methods attempt to capture multiple spike history-dependent effects --- including the recovery period, bursting, and oscillations -- in a single Generalized Linear Model (GLM). To construct the regression model, one considers a long interval of spike history (e.g., spanning (0, 150] ms in the past [12]), divides this interval into time bins of interest, and uses the counts of spikes recorded in each bin as individual predictor variables in the PPM. Typically, a subset of narrow (e.g., 1 ms) bins positioned over the most recent historical timepoints (e.g., -10 through -1 ms) are assumed to capture the combined effects of any short-lag history effects (including the RP and bursting). Wider, longer-lag regressors are positioned to estimate oscillatory effects. For example, spike counts summarized over the -50 to -40 ms bin may be used to estimate 20-25 Hz rhythms [12].

To reiterate, it is useful to recognize how this GLM strategy permits the long-duration ISIs to inform estimates of spike oscillations, rather than discarding this information when segregating out the short-term history effects. In the case of a low-FR unit that spikes rhythmically, but on average only once per oscillation cycle, such a GLM could detect the underlying rhythm, assuming

the cycle duration exceeds the summed width of the short-lag history bins (a reasonable assumption, since most oscillation effects of interest typically fall well below 100 Hz).

The second subcategory consists of a single method, the Latent Oscillatory Spike Train (LOST) model of [11]. LOST is a distinct, Bayesian PPM, which shares some key features with the regression-based models, and also differs in some meaningful ways. Like the above GLMs, LOST preferentially draws upon short-lag spike history effects, as opposed to ISIs of any duration, to capture the recovery period. Instead of using the indicator-function approach of the above GLMs, LOST models the effects of recent spiking as a smooth function of the elapsed time, with this function flexibly captured with a set of cubic splines (with constraints added to pin the estimated function values to 0 for spikes that occurred far in the past). Therefore, LOST achieves the currently prioritized goal of sparing the long inter-spike intervals when controlling for RP effects.

The core difference between the GLMs and LOST entails the modeling of oscillations. LOST represents underlying oscillatory drives as continuous latent states, which are constrained by priors, but may vary over time with respect to their precise modulation strengths and center frequencies. These latent oscillatory components in turn shape the ongoing spike probability.

As emphasized in the Introduction, the chief limitation of these fully time domain-based methods concern (1) a potentially problematic reliance on users' configuration decisions, and (2) the lack of an associated power spectrum estimate. To appreciate the consequences of these limitations, we can first consider the GLM methods. In typical implementations of these methods [12, 13], the bins' temporal offset and width are often set to hard-coded defaults which are applied to all analyzed units. With respect to the goal of modeling the recovery period, this rigid approach neglects cross-neuronal variability in RP duration, the existence of which is indicated by both

analysis of units from our lab (see Chapter 4), and prior research (e.g., [39, 40]). Consequently, the hard-coded short-lag time bins may cover too little time to completely model the RP (resulting in residual distortion effects on the oscillation estimates) or more time than is needed (resulting in unnecessary sacrifice of the ability to estimate the highest frequency oscillations). With respect to the goal of modeling oscillations, the bin-based strategy sacrifices the sharp frequency resolution of a standard DFT-produced spectrum, and represents a departure from the typical basis set of sinusoids. These characteristics may complicate investigations that aim to draw comparisons with previous spectral data, or derive precise oscillation frequency information.

LOST shares these general limitations. For the modeling of the short-lag effects, the user must identify the historical time points at which to place the spline knots, and the bound past which the estimated function will be pinned to zero. Similarly, the modeled oscillatory drives are also chosen *a priori*, and limited to a smaller set of distinct center frequencies than is represented in a typical power spectrum.

In summary, the time-domain methods informed the present work's strategy for isolating the short-lag effects that contribute to the RP, but also left shortcomings to be addressed. Note that both of these shortcomings again touch upon computational themes (e.g., the best choice of model for capturing both oscillatory and short-term aperiodic trends, and the replacement of human judgment with automation). As I will describe in Chapter 3, I will replace the time-domain modeling of oscillations with spectral analysis of each point process model's residuals. Likewise, I will replace the user's configuration of the modeled RP duration with an ISI-driven algorithm that automatically determines this parameter setting.

In addition to addressing the earlier PPMs' limitations, it was also necessary to address a difference in analysis goals. The earlier PPMs indiscriminately lumped together all possible short-

lag effects with a single set of terms, whereas the goal of the primary project proposed here was to selectively account for the recovery period. The recovery period does not refer to the combined effects of all of the most recent spikes (what the prior PPMs modeled), but rather a short-lived effect of the most recent spike only. As I will describe in Chapter 3, I therefore constructed PPMs that selectively modeled the effect of the time elapsed since the most recent spike (when that most recent spike occurred within the sufficiently recent past). This approach was in part informed by another previous method, the Inhomogeneous Markov Interval (IMI) model of [41]. In the IMI model, a set of regression splines modeled the effect of the time elapsed since the most recent spike (in this case, without a bound placed on the historical lag at which the spike could occur in the past). Although the PPMs I propose in Chapter 4 differ from the IMI model in a number of ways, I do adopt this method’s notation for the timestamp of the most recent spike ( $s_*(t)$ ) and for the time elapsed since this most recent spike ( $t - s_*(t)$ ).

## **2.7 Burst Detection: Priorities**

The remainder of this chapter will cover background relevant to the second, exploratory component of this thesis, which aims to account for bursting-related impacts on the power spectrum. To accomplish this objective, it is necessary to first identify which spike trains are indeed “bursty” (i.e., appropriately divided into burst- and non-burst states), and the timestamps when bursts occur. As part of the current work, I developed a new method for this burst detection step. Below, I describe the priorities I set for a prospective burst-labeling method, and in Sections 2.8-2.9 I very briefly review a representative subset of the pre-existing methods.

In my approach to burst detection, I pursued two overarching priorities. First, I sought a detection method that would match the assumptions underlying my ultimate application of the burst labels. Elaborating on the points I made in Introduction, I assumed that:

- (1) “Bursty” spike trains can be reasonably described as alternating between a burst and a non-burst state.
- (2) The two states are distinguished by distinct “baseline firing properties.” At minimum, the states’ mean firing rates (FR) should differ. It is possible that the duration and shape of the recovery period function may also differ.
- (3) Following from (1) and (2), it should be possible to associate each state with distinct baseline firing properties, including a characteristic RP duration (which is critical, to inform the setting of the history window over which the RP is modeled for the two states).

Note that the description of bursters as typically dwelling in two (rather than potentially  $> 2$ ) states is very common in the neuroscience literature (e.g., see [14, 42]). A unit suspected of entering multiple, diverse burst states would need to be analyzed with a method that is more flexible than the one I propose in Chapter 5.

The second priority is a continuation of one I raised in the context of one-state RP correction: I aimed to minimize the demand on the user to configure the free parameters of the burst identification method. I expected that a data-driven procedure would better accommodate neuronal variability, while reducing the influence of potentially arbitrary decisions on the procedure’s output.

## 2.8 Burst Detection: Traditional Methods

In my brief review of existing methods, I will emphasize one report that describes two statistically-grounded, state-inference based approaches to burst detection ([18]; see Section 2.9). Similar to the algorithm I propose in Chapter 5, these prior methods not only segmented spike trains into burst- and non-burst states, but also estimated characteristic baseline firing properties for these two states. However, before describing these state-inference methods, it is important to first acknowledge the vast collection of other burst detectors that have been proposed for spike train analysis (see [14] for a relatively recent review). At odds with my present priorities, these traditional methods do not explicitly aim to learn state-specific firing properties, and in most cases, also rely upon a number of hard-coded parameter settings. However, to provide some context for the state inference methods that follow, I will highlight three representative approaches here:

- (1) The MaxInterval method [43]: This method represents an extreme along the user-defined versus data-driven continuum. The user sets five free parameters (all relating to either the expected within-burst ISIs, or the dwell times within the burst and non-burst states). A hard-coded set of decision rules applies these criteria in order to identify burst segments in the spike train. Methods in this general category (i.e., manually-configured, and rule-based) can be effective if strong *a priori* knowledge exists regarding the typical ISI statistics of the burst and non-burst states, and when these states' ISI features are well-separated in the resulting feature space (e.g., if bursts always exhibit long runs of short ISIs that are never present in the non-burst case). However, such perfect separation is unlikely to hold for many units that neuroscientists seek to study. In cases in which the burst and non-burst ISI features are not well-known in advance, and/or do meaningfully



overlap, methods like MaxInterval are at risk of strong bias towards burst over- or under-estimation, in a manner that depends highly on users' potentially arbitrary decisions.

(2) The LogISI method [44]: LogISI represents a step in the data-driven direction, in that a key parameter setting (the maximum permissible intra-burst  $\log_{10}(\text{ISI})$  value) is determined by an algorithm that is applied to the  $\log_{10}(\text{ISI})$  distribution, rather than being set by the user. However, this algorithm itself relies on other hard-coded free parameters, and a series of hard-coded rules (informed by yet more free parameters) governs the application of the identified maximum  $\log_{10}(\text{ISI})$  to detect the bursts in the spike train. Moreover, the search for the maximum  $\log_{10}(\text{ISI})$  value can fail, in which case the LogISI method reverts to an entirely hard-coded, rule-based approach. Therefore, this procedure maintains much of the risk of biases that are of concern for purely rule-based methods such as MaxInterval.

(3) Poisson Surprise (PS) [45]: The PS method is representative of a family of “surprise”-based techniques [46, 47]. As its core procedure, PS adopts a statistical approach: A burst is detected when the count of spikes encountered in an interval of length  $T$  is deemed sufficiently rare under a null hypothesis of Poisson spiking (with the reference Poisson distribution matched to the unit's overall spike rate, and the threshold for “sufficiently rare” determined by a user-specified cutoff probability). Other methods in the “surprise” family replace the Poisson null model with alternative distributions (a helpful innovation, given the departure from Poisson spiking that the recovery period creates). By comparing local spike counts to a null model parameterized by the unit's global spike rate, PS does represent some progress in addressing the concerns regarding the arbitrary biases that the rule-based models create. However, this method does also rely on a number of hard-coded

parameters (e.g., the probability cutoff, the minimum number of spikes per burst, and multiple values used in a procedure that sets the history window  $T$ ). Therefore, the output of PS can still be highly user-dependent.

## 2.9 Burst Detection: State Inference Methods

The existing publication on state inference-based burst detection [18] presents a variant on standard Hidden Markov Models (HMM) and a Hidden Semi-Markov Model (HSMM). Unfortunately, the specific, proposed methods share the traditional detectors' reliance on many user-defined parameters (in most cases, for initialization of the models' unsupervised training, as described below). However, review of these two methods remains valuable, as these served as an informative precedent for the method that I developed.

This section will emphasize high-level, conceptual descriptions of the proposed HMM and HSMM approaches. Chapter 5 provides a more detailed treatment of the Hidden Markov Model proposed for this thesis (although note that the Chapter 5 approach is simpler than the HMM described below, and closer to the most basic formulation of HMM [48-50]).

### 2.9.1 Recap of Basic HMM Concepts

Before describing the methods from [18], it is useful to recap, in conceptual terms, the standard elements of the basic, discrete-time HMM formulation. Given the observations ( $x$ ) collected over a sequence of timesteps ( $t$ ), an HMM infers the sequence of latent states ( $z$ ) that

gave rise to those observations. The states represent a categorical variable (in our present problem, distributed over the burst and non-burst states). In the most conventional, straightforward HMMs,  $x$  is also a categorical variable, but it is also possible to model  $x$  as a continuous variable [50].

HMMs rely on a Bayesian framework. In the simplest, discrete-time, first-order formulation, an HMM assumes that (1) the first-order Markov property holds, such that the probability distribution over the states  $z(t+1)$  is contingent only upon the state  $z(t)$  (and no prior states), and (2) the distribution over the observations  $x(t)$  is conditioned upon the current state  $z(t)$ , and conditionally-independent of the states and observations at all other timepoints. Therefore, to draw inferences, the HMM relies on three sets of prior probabilities (and note the convention in this thesis of using 1 as the index of the initial timestep):

- (1)  $\pi$ : the probability distribution over the values for  $z(1)$ , the initial state
- (2)  $A$ : the state transition matrix, such that  $a_{ij}$  specifies  $p(z(t+1) = \text{state } j / z(t) = \text{state } i)$
- (3)  $B$ : the observation matrix (for categorical observations) or probability density function (for continuous observations), indicating the probability of observing  $x(t)$  given a latent state  $z(t)$

Equipped with these priors, the HMM may be used to infer the sequence of states  $z(t)$  that maximizes the likelihood of having observed the sequence of  $x(t)$ . In the context of the present problem, the HMM may be used to infer the likelihood-maximizing sequence of burst and non-burst states, given a sequence of observations sampled from the spike train.

A user may either pre-specify the priors, or rely upon unsupervised training of the HMM to learn these priors from a dataset. In the latter case, the unsupervised algorithm still must be initialized, with values set by either the user or some automated process. As also holds for the HMM variant and HSMM discussed below, the choice of initialization values can be

consequential, as the unsupervised algorithm converges to only a local and not necessarily a global maximum [49].

### 2.9.2 Tokdar et al. (2010): HSMM and HMM Approaches to Burst Detection

Tokdar et al. [18] compared a Hidden Semi-Markov Model-based framework for burst detection to a continuous-time HMM variant. The distinguishing features of the HSMM will be explained in the paragraphs below. For the HMM, the continuous-time label refers to a modest difference in the modeling of the state transitions: The application of the  $A$  transition matrix over discrete timesteps (which implies a geometric distribution of dwell times for each state [49]) is replaced with the sampling of dwell times from state-specific exponential distributions ( $f_0^{ITI}$  and  $f_1^{ITI}$ , where ITI = “inter-transition” interval; 0 = non-burst state; 1 = burst state). This approach effectively implements the “memoryless” nature of the Markov property in continuous time, and requires some adjustment to the algorithms applied to the HMM, but the core concepts behind the model remain the same as those underlying the discrete case.

In the Tokdar et al. (2010) HMM, the observations consist of the ISIs of the spike train (which we can subscript with  $k$ , to distinguish the indices of the ISI samples from the clock time  $t$ ). Each sample  $x_k$  is assumed to be drawn from one of two possible state-specific gamma distributions ( $f_0^{ISI}$  and  $f_1^{ISI}$ ; note the gamma assumption is common amongst models that aim to capture the RP and mean FR influences on the ISIs [6]).

Tokdar et al. argued that an HSMM approach might improve upon the state transition assumptions of the continuous-time HMM. By assuming exponential forms for  $f_0^{ITI}$  and  $f_1^{ITI}$ , an HMM maximizes the entropy of these distributions. Tokdar et al. suggested that greater regularity

might characterize real neurons' intra-burst and extra-burst dwell times, and anticipated that replacing the exponential distributions with gamma distributions might aid in better capturing this regularity. The term "semi-Markov" is appropriate for a model that adopts this exponential-to-gamma switch, since the state transitions are no longer memoryless<sup>3</sup> (assuming the gamma shape parameter  $\alpha \neq 1$ ; i.e., the distribution is not equivalent to an exponential distribution [51]).

Tokdar et al. (2010) applied unsupervised training to their HSMM and HMM implementations, to refine the settings for the free parameters (including one parameter for each exponential distribution, and two parameters per gamma distribution). Random values were used for the parameters' initialization. Although a detailed description of the training procedure is beyond the scope of this review, note that the complexity of these models prevented the use of the standard algorithm used in the discrete-time HMM case (i.e., Baum-Welch; see Chapter 5). The authors designed a novel Markov Chain Monte Carlo (MCMC) method to achieve the desired

---

<sup>3</sup>It may be helpful to illustrate the memoryless versus not-memoryless properties of the exponential and gamma distributions in the present context. For example, let the random variable  $X$  represent the dwell time for the non-burst states. Following the notation from [21]: A distribution over  $X$  is memoryless when, for  $t > 0$  and  $h > 0$ ,  $P(X > t+h | X > t) = P(X > h)$ . In other words: Regardless of whether the non-burst state has just been entered, or it has been active for  $> t$  ms, the likelihood of remaining in the state for more than  $h$  additional ms remains the same; that is, the probability lacks any "memory" of the time transpired in the non-burst state. For the exponential distribution, this relationship holds, by virtue of the distribution's flat hazard function. However, for a gamma distribution with  $\alpha \neq 1$ , the hazard function is not flat. If  $\alpha > 0$  (as applies to the gamma distributions used in practice in the Tokdar et al. [18] model), then the hazard function is monotonically increasing [51]. In other words, as time spent in the non-burst state increases, the probability of exiting the state within the next  $h$  ms likewise increases; that is, the probability of state transition is shaped by a "memory" of the time elapsed in the state.

training, but it is important to note that the use of such a Monte Carlo-based method (which, unlike Baum-Welch, involves a random sampling process) does imply that the training outcome is stochastic.

In evaluations on synthetic spike trains, [18] found that the HSMM did outperform the HMM's burst detection performance on some synthetic datasets (e.g., when the simulations both imposed regularity on  $f_0^{ITI}$  and  $f_1^{ITI}$  and introduced sufficient overlap in  $f_0^{ISI}$  and  $f_1^{ISI}$ ), but these methods performed comparably in other cases. However, an arguably more consequential evaluation was reported in the aforementioned review by [14]. In their review, [14] compared eight burst detectors, including the HSMM method, on a large set of synthetic datasets, designed to challenge the detectors in different ways. This analysis revealed a comparatively strong tendency for the HSMM detector to output very high rates of false-positive burst detection, in both never-bursting units (i.e., the simulated ISIs were all drawn from a single Poisson or gamma distribution) and noisy units with a mixture of burst and non-burst spiking.

The above observations raise the question of whether more careful initialization of a state inference-based burst detector might serve to decrease these high false detection rates. In particular, two preparatory steps might prove useful.

First, as I propose in Chapter 5, one might develop an algorithm that judges, based on the observed ISIs, whether a two- or one-state model is likely to better characterize a unit (i.e., whether a unit is “bursty” or not). This need to pre-determine the likely number of hidden states – known as the “order selection” problem -- is a well-recognized challenge in the Hidden Markov Model literature [52]. If a unit is pre-judged to be “non-bursty”, it would be appropriate to drop such a unit from a burst detection analysis.

Second, for the probable “bursty” units, an algorithm could be devised to help inform the initialization of the HMM’s free parameters, prior to refinement of these parameters by an unsupervised training method. Such data-driven initialization (which I propose in Chapter 5) may help mitigate the risk that the unsupervised training method terminates on a highly suboptimal local solution. Note that the assumption that parameter initialization decisions can meaningfully impact burst-labeling outcomes can be verified by experimenting with the initialization settings in the HSMM and HMM code supplied by [18] (see <https://www2.stat.duke.edu/~st118/Software/>).

### **3.0 One-State Residuals Method: Materials and Methods**

#### **3.1 Publishing Note**

The text and figures in this Chapter have been adapted from a manuscript that has been posted to a preprint server [19] and submitted to a peer-reviewed journal. The text in Subsections 3.5.3-3.5.5 was primarily contributed by one of the manuscript co-authors.

#### **3.2 Ethics Statement**

The single nonhuman primate (NHP) dataset analyzed in this study (see “Experimental Data Methods”) was acquired under a protocol approved by the Institutional Animal Care and Use Committee of the University of Pittsburgh (#18093682). Experimental procedures and animal care complied with the National Institutes of Health Guide for the Care and Use of Laboratory Animals, the PHS Policy on the Humane Care and Use of Laboratory Animals, and the recommendations of the Weatherall report. The animal was housed with a single cage mate in a climate-controlled room, received regular enrichments, and consumed a diet that included fresh fruits and vegetables daily. Surgical procedures were conducted under general anesthesia, with analgesics administered post-operatively.



### 3.3 Methods Overview

In the subsections that follow, we first describe the procedures for comparing the two correction methods (shuffling and residuals) on synthetic spike trains. These include the steps for generating the spike trains, calculating the original and corrected PSDs, and evaluating the methods' accuracy in detecting known oscillations. Subsequently, we describe the experimental procedures used in acquiring the NHP data, and also the few details of PSD analysis that differed between the NHP data and the synthetic spike trains.

All simulations and analyses were implemented in Matlab (R2022a, The Mathworks, Inc., Natwick, MA; RRID:SCR\_001622). Matlab code sufficient to reproduce the results reported in the main text of this thesis is available on Github ([https://github.com/kc13/residuals\\_spectral\\_analysis](https://github.com/kc13/residuals_spectral_analysis); DOI: 10.5281/zenodo.10867519); code for the results reported in the Appendices is available upon request. We make use of conventional Matlab notation for referring to the ranges spanned by / increments separating vector elements ([min:step:max]).

Data files are available in both a Zenodo repository (DOI: 10.5281/zenodo.8313070) and the Github repository. The original synthetic datasets and NHP spike train data are available on Zenodo. The Github repository stores copies of the NHP spike train data, and also synthetic datasets that are represented in a postprocessed state (the subsampling and partial ROC output, see “Comparison of ROC curves for the Correction Methods”; note the original synthetic datasets exceed Github file size limits).

## 3.4 Synthetic Data Methods

### 3.4.1 Spike Train Simulation

We generated synthetic spike trains through the procedure described in [7]. For each time bin  $t$  ( $\Delta t = 1$  ms), the spike probability  $p_{spk}(t)$  was determined according to the discrete-time approximation of an inhomogeneous Poisson process:

$$p_{spk}(t) = \begin{cases} p_{SS}(t) = p_{base} + p_{osc} \cdot \sin(2\pi f_{osc} t \cdot 0.001), & n > n_r \\ p_{RP}(t) = p_{SS}(t) \cdot k^{n_r+1-n}, & n \leq n_r \end{cases} \quad (1)$$

Here,  $n$  indexes the millisecond latency since the last spike occurrence, and  $n_r$  denotes the duration of the recovery period (RP). Therefore, the spike probability at  $t$  depends on whether the simulated unit is in the recovery period (RP) or has recovered the steady state (SS) firing pattern. During the steady state,  $p_{spk}(t)$  reflects the sum of two components: (1)  $p_{base}$ , corresponding to the mean steady state firing rate (FR), and (2) an oscillation term, in which  $0 \leq p_{osc} \leq p_{base}$  governs the oscillation amplitude, and  $f_{osc}$  (specified in Hz) indicates the oscillation frequency. Following from [2], we set  $p_{osc}$  indirectly through variation of the modulation index,  $m = p_{osc}/p_{base}$ ; i.e., the ratio of the peak and mean steady state FRs, respectively. The above  $p_{osc}$  constraint implies that  $0 \leq m \leq 1$ . For all simulations,  $p_{base} \ll .5$ , thereby implying that  $p_{spk}(t)$  remained bounded in  $[0,1]$ .

During the recovery period,  $p_{spk}(t)$  is set equal to the steady state probability following modulation by an exponential recovery term, where  $0 \leq k < 1$ . Following from [7], our simulations adopted default RP parameter settings of  $k = 0.7$  and  $n_r = 9$ ; a set of follow-up tests considered the use of other values (see ‘‘Evaluation of the PSD Correction Methods’’).

The duration of the synthetic spike trains ( $T$ ) was constrained to multiples of 1024 ms (the segment size used for spectral analysis). The settings used for  $T$  and three other free parameters ( $p_{base}$ ,  $m$ ,  $f_{osc}$ ) are described in further detail in “Evaluation of the PSD Correction Methods”. The output of each simulation run was a  $T \times 1$  binary vector (a “delta vector”), indicating time bins of spike occurrence at a resolution of 1 kHz.

We executed a small number of additional simulation runs, using a different simulation framework. These additional simulations addressed questions concerning method performance on spike trains generated to express certain biologically-realistic features. Appendix A (section A.2.1) describes the methods for these additional simulations.

### 3.4.2 Uncorrected PSD

To obtain the initial, uncorrected PSD estimates, we submitted the delta vectors (demeaned segment-wise) to Welch’s method (Matlab *pwelch*, Hamming window length =  $NFFT = 1024$  ms,  $Fs = 1$  kHz, zero segment overlap). PSD functions were represented as 513-element vectors spanning the frequency range [0, 500] Hz, implying a frequency resolution of  $\Delta f \approx 0.9766$  Hz.

Within each PSD function, we sought points of statistically significant power in the range (0,100] Hz. Significance was determined using the CI-construction procedure from [7] which was described in the Introduction. For our analyses, points in the (0, 100] Hz range of a spectrum were labeled as significant if they exceeded the upper limit of a one-sided,  $[100*(1-\alpha_c)]\%$  confidence interval formed using the mean and SD over the [250, 500] Hz range of the same spectrum ( $\alpha_c =$  Bonferroni-corrected significance level, with respect to the 102 points in the (0,100] range). All figures of uncorrected PSDs were generated with  $\alpha_c = .05$ .

For the shuffling and residuals methods detailed below, all spectral analysis and statistical testing steps made use of the same routines as were applied to the uncorrected PSDs, and the same default  $\alpha_c = .05$  for figure generation;  $\alpha_c$  was varied away from this default for the ROC analyses described in “Evaluation of the PSD Correction Methods”.

### **3.4.3 Shuffling-corrected PSD**

To obtain the shuffling-corrected PSD for a given spike train, we first generated 100 surrogate, ISI-shuffled versions of the original delta vector (Fig. 2B). We implemented the global shuffling method described by [7] (i.e., random permutation of all ISIs across the entire vector). We chose this approach over local shuffling (permutation of ISIs within nonoverlapping segments) since our simulations emphasized relatively low firing rates (necessitating a broad shuffling scope to admit a sufficient number of ISIs) and did not use time-varying simulation parameters (eliminating the need to capture local variations). Additionally, using the same primary evaluation that we describe for the shuffling and residuals methods (see “Evaluation of the PSD Correction Methods”), we verified superior performance for global as compared to local shuffling. Code sufficient to reproduce this observation is available in the Github repository.

Each of the 100 globally shuffled spike trains was submitted to spectral analysis, and the mean of the resulting PSDs was divided from the original uncorrected PSD, producing the final shuffling-corrected PSD.

### 3.4.4 Residuals-corrected PSD

The residuals method entails three steps: (1) estimation of the recovery period duration, (2) fitting a point process model (PPM) to the spike train, which estimates the effect of the time elapsed since the most recent spike, up to a history bound corresponding to the estimated RP duration, and (3) applying spectral analysis to the residuals of the PPM fit.

#### 3.4.4.1 Estimation of Recovery Period Duration

We define the recovery period as a post-spike reduction in a unit's spike probability, terminating upon return to a steady state firing rate. For our synthetic spike trains (excluding a series of follow-up datasets), the ground truth duration of the recovery period was  $n_r = 9$  ms.

Our estimation approach builds upon a general logic that was described in [53], although note that these authors did not propose an explicit estimation algorithm. We expand upon these authors' ideas by detailing a specific algorithm that operates on the probability density function (PDF) of the inter-spike intervals (ISI). To introduce our approach, we use a simple synthetic spike train for which  $m = 0$  (i.e., no oscillation present, such that firing is shaped by only  $p_{\text{base}}$  and the RP); however, subsequent analyses will show that the algorithm performs reasonably when  $m > 0$ .

The ISI PDF for an example spike train generated with  $m = 0$  is shown in Figure A5.A. Recall that, in the case of Poisson firing (i.e.,  $n_r = 0$ , with an intensity  $\lambda = p_{\text{base}}$ ), the observed ISIs are sampled from an exponential PDF [6]:  $P(\text{ISI} = x) = \lambda e^{-\lambda x}$ . Because a spike train with an RP is not Poisson, an exponential curve fit to such a spike train's ISI PDF will overestimate the probability mass assigned to the ISI lags  $\leq n_r$ . Now, consider what we should observe if we progressively shift the exponential curve fit rightward. Figure A5.5A illustrates this process for a series of example

lags  $L$  assigned to the  $x = 1$  position of the exponential fit. When  $L \leq n_r$ , the exponential curve continues to overestimate the probability mass for the lowest lags. However, when  $L = n_r + 1$ , the exponential model of this left-cropped ISI distribution is appropriate, as firing outside of the RP is Poisson. Our strategy for identifying the boundary between the end of the RP and the start of the steady state relies on this expected transition. We additionally leverage the observation that statistical evidence in favor of the exponential fit (i.e., the extent to which the addition of an exponential curve explains variance beyond that explained by a constant term alone) tends to be stronger when  $L = n_r + 1$  than when it is incremented to the immediately neighboring lags to the right.

The estimation algorithm proceeded as follows:

1. Given a delta vector, compute the ISI histogram, with bins spanning the ms lags in  $[1, \max(\text{ISI})]$ .
2. For each  $L$  in *lags*:
  - a. Extract the histogram values over the range  $[L, \max(\text{ISI})]$ .
  - b. Normalize the extracted histogram to unit area to form a PDF.
  - c. Fit a single scaled exponential to the PDF using the generalized linear model (GLM) of the form  $PDF(x) = e^{B_0} e^{B_1 x} \rightarrow \log(PDF(x)) = B_0 + B_1 x$ . See the next subsection for more detail on this form of GLM.
  - d. Compute the deviance difference, which is a standard statistic for assessing the improvement in goodness-of-fit contributed by a full GLM (“model 1”, here set to the constant term plus the exponential lag effect), compared to a GLM restricted to the constant  $B_0$  term (“model 0”). The difference statistic is  $\Delta D = D_0 - D_1 = -2[\log \hat{L}_0 - \log \hat{L}_1]$ ; i.e., the log likelihood ratio scaled by -2 [6, 54].

- e. Upon confirming the first instance of a  $\Delta D$  value that is a local maximum ( $\Delta D(L-1) < \Delta D(L) > \Delta D(L+1)$ ), break out of the loop iterating over *lags*.
3. The  $L$  associated with the local  $\Delta D$  maximum marks the start of the steady state region; the estimated recovery period duration,  $\hat{n}_r$ , is set to  $L-1$  ms.

For the synthetic unit with  $m = 0$ , this process of finding the lag  $L$  associated with the local  $\Delta D$  maximum is illustrated in Figure A5.B. Note that in this case, the duration of the recovery period is accurately estimated as  $L-1 = 9$  ms.

Figure 3A-B illustrates the same process as applied to the two synthetic oscillatory spike trains introduced in Figure 2. Although the presence of non-zero oscillation (i.e.,  $m > 0$ ) implies that the steady-state firing is not Poisson, we observe that the curve-fitting algorithm is robust to this departure and again correctly identifies the ground truth RP duration of 9 ms. Evaluation of the RP estimation accuracy over a full dataset of synthetic spike trains is presented in the Results.

#### 3.4.4.2 Bounded Last-Spike Point Process Model

Once the RP duration is estimated, the next step is to estimate the RP shape. For this purpose, we fit the spike train with a point process model (PPM).

With reference to [6, 8, 11, 41, 55], we will briefly review the fundamental concepts and terminology for point process models in general, and for our implemented Poisson GLM in particular. A PPM seeks to predict the likelihood that a discrete event (e.g., a spike) will occur at time  $t$ , given a set of covariates. For a stochastic point process represented in continuous time (e.g., a spike train expressed as a series of spike timestamps), the PPM models the conditional intensity function (CIF):

$$\lambda[t|\theta, H(t)] = \lim_{\Delta t \rightarrow 0} \frac{P[N(t + \Delta t) - N(t) = 1 | \theta, H(t)]}{\Delta t} \quad (2)$$

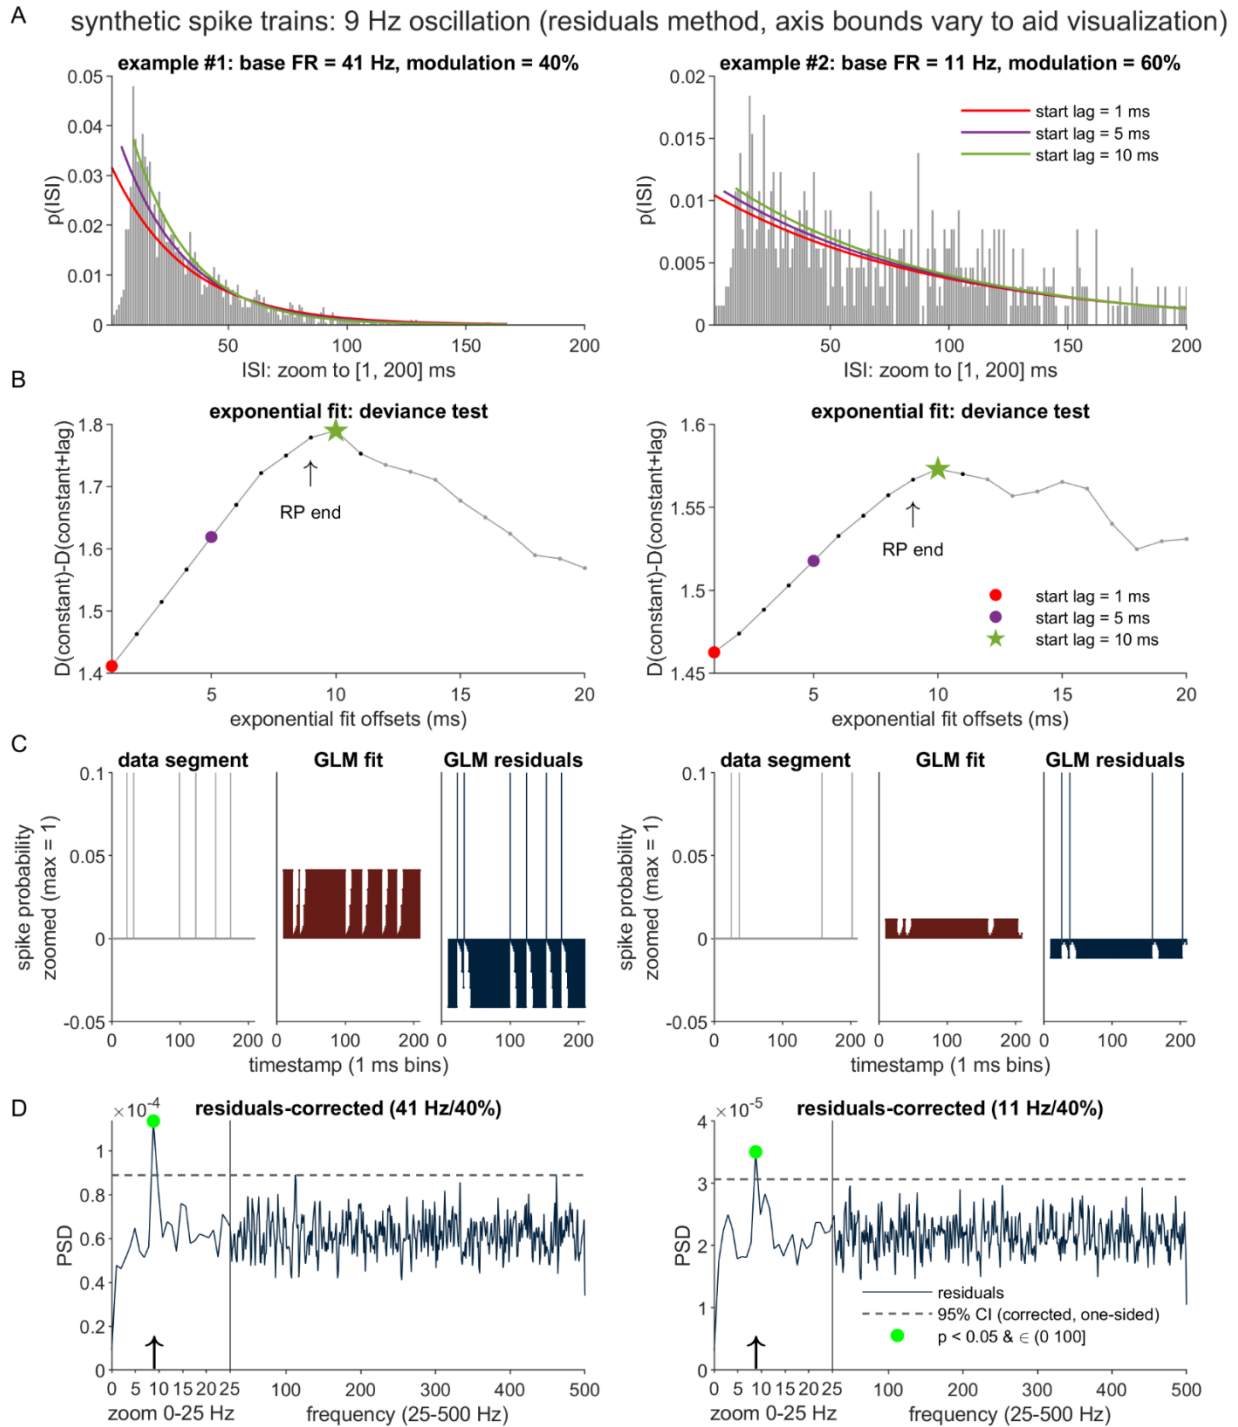
where  $P[N(t + \Delta t) - N(t) = 1 | \theta, H(t)]$  is the conditional probability of a discrete event instance (e.g., a spike) occurring in the time interval  $(t, t + \Delta t]$ , given a covariate history  $H(t)$  and model parameters  $\theta$ . For our discrete-time spike train representation (the delta vectors), this continuous formulation may be approximated using the assumption that for  $\Delta t \leq 1$  ms,  $P(\text{spike in } (t, t + \Delta t]) \approx \lambda[\theta, H(t)]\Delta t$ . For the remainder of this text, as a shorthand, we will follow the convention of using  $t$  to index the ms-width bins of the discrete-time representation (as opposed to continuous time).

Given this discrete-time representation, a Poisson GLM (i.e., a generalized linear model with a log link function, also known as Poisson regression) is one example of a PPM that may be used to predict spike probability. We implemented the bounded last-spike model as a Poisson GLM of the form

$$\log \lambda(t|H_t, \beta) = \beta_0 + \sum_{j=1}^{\hat{n}_r} \beta_j \delta_{t-j, s_*(t)} \quad (3)$$

where  $s_*(t)$  denotes the time of the most recent spike [41], and  $\delta_{t-j, s_*(t)}$  is the Kronecker delta function, returning 1 when  $t - j$  equals  $s_*(t)$ , and 0 otherwise. In other words, the model estimates the variance in the logarithm of spike intensity that is explained by either the constant term alone (outside of the RP) or this same constant plus the weighted contribution of the millisecond interval transpired since the last spike (inside the RP). Consequently, inside the RP,  $\lambda$  is modeled as the





**Figure 3. Distortion correction by the proposed residuals method.**(A) Illustration of the procedure for obtaining an estimate of the RP duration ( $\hat{n}_r$ ), as applied to the two example spike trains from Figure 2. A series of right-shifted exponential curves are fit to each ISI distribution, left-anchored to starting positions advanced in 1 ms steps (with 3 sample iterations highlighted in the figures). (B) Plots of the deviance difference statistic,  $\Delta D$ , as a function of the first 20 starting positions of the exponential fits.  $D(\text{constant})$ ,  $D(\text{constant}+\text{lag})$  = deviance measures for the intercept-only and intercept+exponential curve models, respectively.  $\Delta D$

tracks the goodness of fit contributed by the exponential curve. Each  $\hat{n}_r$  estimate is set equal to the post-spike lag immediately preceding the first local maximum in the corresponding  $\Delta D$  plot. (C) Illustration of the Poisson Generalized Linear Model (GLM) accounting for the effects on spike likelihood of the time elapsed since the most recent of any spikes occurring within the preceding  $\hat{n}_r$  ms. The three panels (which highlight the first 209 ms of data) depict the GLM input (left panels), model fit (middle panels), and the model residuals (right panels). (D) PSDs of the GLM residuals. Plotting and statistical conventions follow those from Figure 2.

multiplicative contribution of the exponentiated constant and last-spike factors ( $e^{\beta_0} e^{\beta_j}$ , for  $j = t - s_*(t)$ ), and as  $e^{\beta_0}$  outside of the RP<sup>4</sup>. The model was fit to the spike train data using Matlab *fitglm*, which applies the conventional iteratively re-weighted least squares (IRLS) method to seek the log likelihood-maximizing parameters. Note that the log likelihood function for a model such as (3) is known to be concave [8], thereby implying that if a maximum exists, it must be unique.

#### 3.4.4.4 Spectral Analysis of Model Residuals

Given a set of parameter estimates  $\hat{\beta}$ , the raw residuals of the modeled spike train (the observed - fitted values) are calculated as

$$r(t) = I(t) - e^{\left[\hat{\beta}_0 + \sum_{j=1}^{\hat{n}_r} \hat{\beta}_j \delta_{t-j, s_*(t)}\right]} \quad (4)$$

where  $I(t)$  represents the binary spike train input to the PPM. The final step of the residuals method is submission of  $r(t)$  to spectral analysis. Note that (4) only defines  $r(t)$  for  $t \geq \hat{n}_r + 1$  ms. To

---

<sup>4</sup> As a concrete example: Consider a case in which  $n_r = 9$  ms, and relative to time  $t = 0$ , the last two spikes occurred at  $t = -2$  and  $t = -4$ . Because the most recent spike occurred 2 ms in the past, timepoint  $t$  falls in the RP, and spike intensity is predicted by  $e^{\beta_0} e^{\beta_2}$ .

maintain alignment with the original time series (and the divisibility by the 1024 ms segment size), the initial 1024 ms segment submitted to *pwelch* consisted of the initial 1024- $\hat{n}_r$  ms of computed residuals, demeaned and prepended with zero-padding of  $\hat{n}_r$  ms.

### 3.4.5 Evaluation of the PSD Correction Methods

Comparison of the residuals and shuffling methods entailed (1) simulation of synthetic spike trains over a range of physiologically-relevant parameters, (2) generation of the two methods' corrected PSDs, with tabulation of hit and false alarm rates for these PSDs across a range of  $\alpha_c$  levels, and (3) comparison of the methods' classification of true and false oscillations, by way of ROC (Receiver Operating Characteristic) analysis (collapsed across parameter settings) and regression analysis of the residuals-shuffling differences in the hit and false alarm rates (to examine parameter effects).

#### 3.4.5.1 Synthetic Spike Train Test Sets

Method evaluation was first carried out on a primary synthetic test set; subsequently, we generated a number of additional test sets to address follow-up questions. For each test set, and for each of the unique combinations resulting from the crossing of the varied simulation parameters, we randomly generated one hundred spike trains.

For the primary test set, 100 spike trains were simulated for each of the 540 unique combinations that resulted from crossing the following four varied parameters:

- (1) Spike train duration ( $T$ ):  $\{30,60,120\} \times 1024$  ms windows
- (2) Oscillation frequency ( $f_{osc}$ ):  $\{7,9,12,20,32\}$  Hz

(3) Base firing rate offset<sup>5</sup> ( $p_{base\_offset}$ , such that  $p_{base} = f_{osc} + p_{base\_offset}$ ):  $+2^{[0:1:5]}$  Hz

(4) Modulation index ( $m = p_{osc}/p_{base}$ ): [0:0.2:1]

The above parameters emphasize relatively low duration and firing rate settings, with the short durations consistent with the limitations that can be posed by real experimental data, and the low FRs in line with the motivation for developing the residuals method. At the same time, the setting of  $p_{base}$  as a positive offset from  $f_{osc}$  ensured that, in most cases, the synthetic unit averaged  $> 1$  spike per oscillation cycle, although the recovery period implied that the average dropped to slightly  $< 1$  in some of the  $f_{osc} + 1$  cases. For  $m$ , note that the inclusion of 0 allowed for tracking of false positive detections of significant power when no oscillatory effect was introduced into the simulation.

Three kinds of variants on this primary dataset were created to address follow-up questions. First, we generated one dataset in which  $p_{base}$  and  $p_{base\_offset}$  were directly manipulated, leaving  $f_{osc} = p_{base} - p_{base\_offset}$  as the indirectly manipulated variable. This dataset allowed for tests to determine whether the methods' relative performance was significantly influenced by the magnitude of the  $p_{base} - f_{osc}$  difference, when controlling for  $p_{base}$  itself. For this dataset, the simulations used the same  $T$  and  $m$  settings as were used for the primary dataset. We crossed these factors with two levels of  $p_{base}$  (15 Hz, 20 Hz) and 4 levels of  $p_{base\_offset}$  ([2:2:8] Hz), resulting in 144 unique parameter combinations.

---

<sup>5</sup> As clarification: This relationship assumes that  $p_{base}$  and  $p_{base\_offset}$  can simultaneously refer to terms that contribute to the probability of spiking within a 1 ms time bin, or to the expected number of spikes per second (Hz). This assumption does imply that the value of these terms changes depending on the unit under consideration (e.g.,  $p_{base} = 11$  Hz expected spike rate, or a .011 probability of spiking in a bin).

The second and third variations involved the use of higher FRs, and alternatives to the default RP parameters, respectively. For the high-FR simulations, the set of sampled  $p_{base\_offset}$  values was changed to [35:10:85] Hz; these were added to the  $f_{osc}$  values used for the primary dataset ((2) in the list above) to determine the  $p_{base}$  settings. For the altered-RP simulations, we tested both shorter and longer RP durations ( $n_r$ ) and varying exponential steepness terms ( $k$ ); full details are described in the Results.

### 3.4.5.2 Labeling of Hits and False Alarms

For each of the 540 x 100 (or 144 x 100) synthetic spike trains in a test set, the shuffling and residuals methods were both applied to produce corrected PSDs. Given a corrected PSD and an  $\alpha_c$  setting, statistical testing returned a binary vector indicating whether oscillatory power exceeded the CI bound for each of the 102 frequency bins of width  $\Delta f = 0.9766$  Hz in (0,100] Hz. Frequency bins are labeled by their lower bound (0.9766, 1.9531, 2.9297...). Let  $f$  denote the frequency labels and  $s(f)$  denote the binary significance output. We classified a PSD as containing a “hit” (i.e., a true positive detection of oscillatory power) if (1)  $m > 0$  and (2) any  $s(f) = 1$  for those labels  $f_i \in f$  that were the 3 nearest neighbors of  $f_{osc}$  (e.g.,  $f_i \in \{10.7422, 11.7188, 12.6953\}$  for  $f_{osc} = 12$  Hz). We classified a PSD as containing a false alarm (FA) if either of two cases held true: (1)  $m = 0$ ,  $s(f) = 1$  for any label  $f_i \in f$ , or (2)  $m > 0$ ,  $s(f) = 1$  for  $f_i \in |f_i - f_{osc}| > 5$  Hz (equivalent to those  $f_i$  outside the 10 nearest neighbors of  $f_{osc}$ ).

Note that the output of this classification procedure was a pair of binary labels for each PSD, indicating whether each contained  $\geq 1$  hit, and  $\geq 1$  FA, respectively. For the ROC analysis described below, counts of hits and FAs refer to the counts of these PSD-level labels (as opposed to the total counts of individual hit- or FA-classified frequency bins).

### 3.4.5.3 Comparison of ROC curves for the Correction Methods

ROC analysis of the shuffling and residuals methods consisted of three steps: (1) random selection of subsamples of the PSDs, to be used for a bootstrapping procedure for statistical testing, (2) calculation of partial ROC curves and corresponding partial area under the curve (pAUC) values for the subsampled PSDs, and (3) statistical testing of the mean difference in the pAUC values yielded by the shuffling and residuals methods. Note that the ROC analyses were limited to the  $m > 0$  cases, thereby implying that, for those datasets that originally consisted of 540 unique parameter conditions, 450 (unique parameter conditions)  $\times$  100 ( $n$  spike trains per condition) were available for ROC analysis.

Resampling procedures are commonly used as a means of estimating confidence intervals in the context of ROC analysis [56]. As an overview, we adopted a procedure of repeatedly drawing small subsamples of corrected PSDs (20% of the dataset), and deriving the shuffling- and residuals-generated pROC curves for each of those subsamples (see Figure 4D for an example). The outputs of this step enabled estimation of the variance in the average difference in the methods' pAUC values, thereby allowing statistical comparison of this quantity against 0.

The subsampling procedure was iterated 1000 times. For each subsampling iteration, we randomly sampled, without replacement, 20 of the 100 synthetic spike trains -- and the corresponding shuffling- and residuals-corrected PSDs -- generated within each of the 450 parameter combinations. Using the resulting subsample of  $450 \times 20 = 9000$  pairs of shuffling and residuals PSDs, and the classifications of these PSDs as hit and/or FA-containing, we tabulated the two methods' respective hit and false alarm rates over all of the 9000 samples (thereby collapsing over the represented parameter combinations). These hit and false alarm rates could

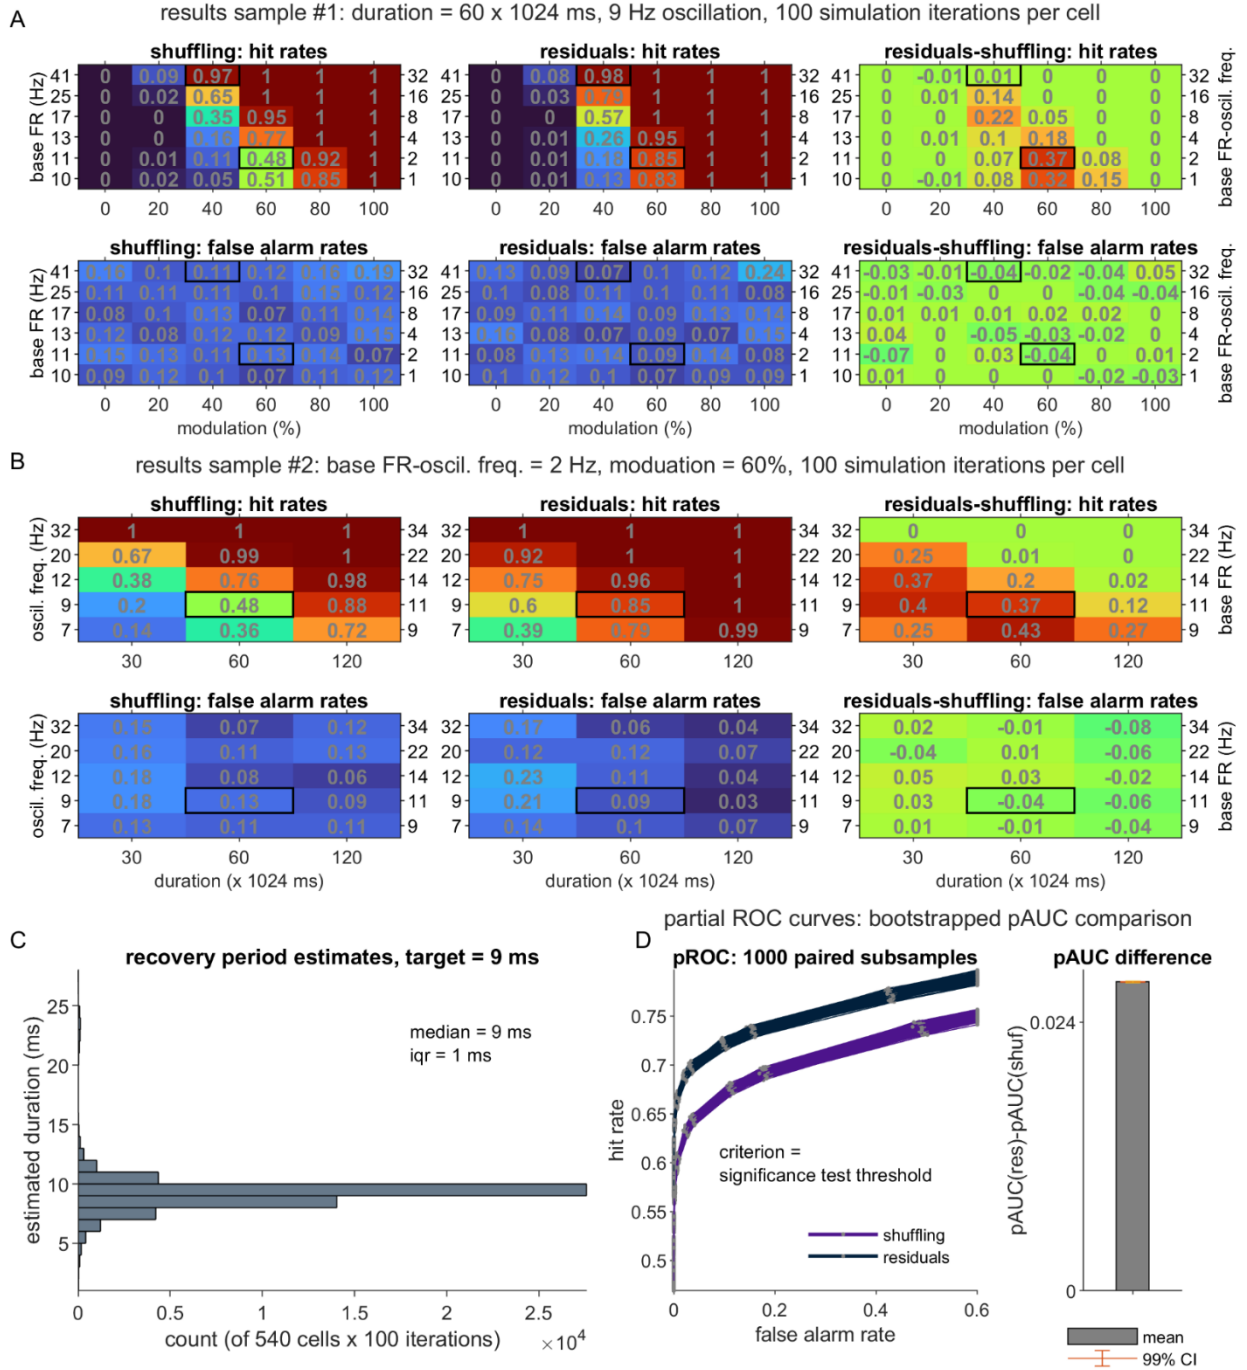
then be used to construct the methods' ROC curves corresponding to that particular subsampling iteration.

For a given subsampling iteration and correction method, we formed an ROC curve as the function relating hits to FAs over a range of  $\alpha_c$  values. The full set of  $\alpha_c$  values consisted of the sorted vector of the values 1 and 5 multiplied by a series of negative powers of 10 ( $[1 \ 5]^T * [10^{-8}:1:-1]$ ) and 1. Therefore, a single point on an ROC curve represented, for one correction method, the fraction of the 9000 corrected PSDs that produced  $\geq 1$  hit and  $\geq 1$  FA, respectively, at a given  $\alpha_c$  level.

For each analyzed dataset (i.e., the primary dataset, the high-FR dataset, and the three altered-RP datasets), the above procedure resulted in 1000 pairs of shuffling- and residuals-generated ROC curves. Note that the FA and HR ranges for these curves could vary slightly, and fall short of spanning  $[0,1]$ . Curves that did not reach 100% HR or FA rates (even at  $\alpha_c = 1$ ) were possible due to the use of a one-sided significance test: If all points in either the hit or FA zones of a PSD were numerically less than the mean of the high frequencies-based CI, then these points could never pass even the most lax significance threshold. Curves that did not drop to 0% HR or FA rates reflected PSDs with prominent hit- or FA-eligible points that survived even a  $\alpha_c = 1 \times 10^{-8}$  threshold (which is much more stringent than the criteria that would be typically applied in practice; e.g., [7]).

The variability in the curves' spanned ranges precluded analyses that compared the areas under the entirety of the ROC curves. Therefore, we instead adopted the partial area under the curve (pAUC) approach [57], such that we exclusively considered the area under an ROC region that was shared by all 2000 curves. Specifically, the raw curves were cropped (using linear interpolation as needed) to span the intersection of all 2000 curves' FA ranges.

synthetic spike trains: method evaluation (see caption for hit/FA definition; "hit" is FA if modulation = 0%)



**Figure 4. Performance of the shuffling and residuals methods over a diverse synthetic dataset.** One hundred synthetic spike trains were generated for each of 540 unique parameter setting combinations (resulting from the crossing of the varied duration ( $T$ ), modulation strength ( $m$ ), oscillation frequency ( $f_{osc}$ ), and [base FR - oscillation frequency] offset ( $p_{base\_offset}$ ) parameters). For each spike train, corrected PSDs were generated by the two correction



methods. When  $m > 0$ , a PSD represented a “hit” if it contained  $\geq 1$  point of significant power within the three nearest neighbors of the ground truth oscillation frequency ( $= f_{osc}$ , when  $m > 0$ ) on the sampled frequency axis (resolution = 0.9766 Hz). A false alarm (FA)-containing PSD included  $\geq 1$  point of significant power at frequency points beyond the 10 nearest neighbors of  $f_{osc}$  (when  $m > 0$ ) or at any frequency point (when  $m = 0$ ). (A) A subset of the methods’ hit and FA rates, shown for all combinations of  $m$  and  $p_{base\_offset}$ , for a fixed  $T = 60 \times 1024$  ms and  $f_{osc} = 9$  Hz. Bold outlines highlight the parameter combinations used for the two example spike trains depicted in Figure 2-Figure 3. (B) Hit and FA rates, shown for all combinations of  $T$  and  $f_{osc}$ , for a fixed  $p_{base\_offset} = 2$  Hz and  $m = 60\%$ . Bold outlines highlight the parameter combinations used for the second, lower FR spike train depicted in Figure 2-Figure 3. (C) Histogram of the residuals method’s estimates of the recovery period duration (ground truth  $n_r = 9$  ms), collapsed over all of the synthetic spike trains. (D) Partial receiver operator characteristic (pROC) analysis of the residuals (*res*) and shuffling (*shuf*) methods’ performance, restricted to the 450  $m > 0$  conditions (for which a hit rate may be defined). Left: Hit versus FA rates are plotted over varying statistical test thresholds; individual curves represent spike train subsamples (matched for *res* and *shuf*) that were uniformly drawn from the 450 conditions (to obtain bootstrapped estimates of variability). The “partial” designation refers to the truncating of individual curves to the FA range bounded by the minimum and maximum FAs observed over all curves. Individual curves may fail to reach 100% hit and FA rates due to the use of a one-sided statistical test, and may remain above 0% due to highly elevated PSD values that surpass even the most stringent  $\alpha_c$  level ( $1 \times 10^{-8}$ ). Right: Mean and 99% CI corresponding to the  $t$ -test ( $df = 999$ ) contrasting the residuals-shuffling partial area under the curve (pAUC) difference against zero.

For each cropped curve, we used the trapezoidal rule (Matlab *trapz*) to compute the area under this bounded curve region. Due to the subsampling process, the paired differences in the shuffling and residuals methods’ pAUC estimates ( $D_{pAUC}(i) = pAUC_{res}(i) - pAUC_{shuf}(i)$ ) were assumed to approximate a normal distribution; therefore, a two-tailed paired  $t$ -test was used to assess the null hypothesis of  $D_{pAUC} = 0$ .

#### 3.4.5.4 Analysis of Simulation Parameter Effects on Relative Hit and False Alarm Rates

To examine the simulation parameter effects, we focused on a fixed  $\alpha_c = .05$  threshold, and applied regression analyses to the residuals-shuffling differences in HRs ( $D_{HR}$ ) and FAs ( $D_{FA}$ ) separately. To avoid specific distributional assumptions regarding  $D_{HR}$  and  $D_{FA}$ , we conducted these analyses using the same 1000 subsamples as were used for the pROC tests. We considered all simulation parameter combinations, with the exception of those for which  $m = 0$  (to omit cases when  $D_{HR} = 0 - 0$  by definition, and with this same exception applied to the  $D_{FA}$  analyses for consistency). Consequently, 450 x 1000 (or 120 x 1000) unique  $D_{HR}$  and  $D_{FA}$  estimates, respectively, contributed to the two analyses, for each dataset on which they were performed.

Separate general linear models for the two dependent variables (DVs),  $D_{HR}$  and  $D_{FA}$ , were constructed using five independent variables (IVs): the mean-centered equivalents of the four directly manipulated simulation parameters ( $T$ ,  $f_{osc}$ ,  $p_{base\_offset}$ , and  $m$  for most datasets, with  $f_{osc}$  replaced with  $p_{base}$  when  $p_{base}$  was directly manipulated), and also the square of the mean-centered modulation strength parameter ( $(m-\bar{m})^2$ ), given trends in the data visualizations that were suggestive of quadratic effects of the  $m$  factor (Figure 4A-B, Figure A2). For each DV, the full GLM consisted of the intercept term plus the sum of the main effect, 2-way interaction, and 3-way interaction terms formed from the five IVs.

#### 3.4.5.5 Post-Hoc Analyses of False Alarms in Spike Trains with High Firing Rates

In addition to the above ROC and regression analyses, we applied an additional set of post-hoc analyses to the synthetic dataset that had been generated with high FR settings (see “Synthetic Spike Train Test Sets” above). These follow-up assessments were motivated by the observation of

unexpectedly high FA rates, especially for spike trains generated with long  $T$  and high  $m$  settings, and either low  $f_{osc}$  (residuals method) or high  $f_{osc}$  (shuffling method; see Fig. A9-A11) settings. We report three analyses, all of which summarized over the 1000 spike train subsamples, and focused on the  $T = 120 \times 1024$  ms and  $m = 1.0$  conditions.

The first analysis was a linear regression consisting of two IVs ( $f_{osc}$ ,  $p_{base\_offset}$ ) and the DV (the subsampled  $D_{FA}$  scores). Consequently, 30 (= number of  $f_{osc} \times p_{base\_offset}$  combinations)  $\times$  1000 (=  $n$  subsamples) unique  $D_{FA}$  estimates contributed to this analysis. The regression model included both main effects and the  $f_{osc} \times p_{base\_offset}$  interaction term.

The second and third analyses entailed the calculation of descriptive statistics. These statistics were computed separately for the two correction methods. For one set of descriptive statistics, we report the mean absolute FA rate for each of the two methods and the two most extreme  $f_{osc}$  cases (7 Hz, 32 Hz); this value simply corresponded to the grand average of the FA rates, collapsed over the 1000 subsamples and six  $p_{base\_offset}$  settings. For the other set, we report summaries of the peak power (and associated frequencies) for the PSDs' false alarm points, expressed as a ratio over the peak power for the hit points. For each correction method, we iterated over the 1000 subsamples of spike trains, and for each subsample, over the 30 unique  $f_{osc} \times p_{base\_offset}$  setting combinations. Within each subsample, each unique  $f_{osc} \times p_{base\_offset}$  setting combination is represented by 20 spike trains (given the constraint that  $T = 120 \times 1024$  ms and  $m = 1.0$ ). For the correction method under consideration, we identified the FA and hit points present within the 20 respective corrected PSDs. For each PSD, the peak (i.e., maximum) power estimates were identified over any identified FA and hit points, respectively. For a PSD that included both FA and hit points (as was common in the  $T = 120 \times 1024$  ms,  $m = 1.0$  cases), we computed the  $peak_{FA} / peak_{Hit}$  power ratio. These peak ratios and the associated peak FA frequencies were then averaged over the PSDs (for which they

were defined), and then averaged over the 30 parameter setting combinations and 1000 subsamples.

## **3.5 Experimental Data Methods**

### **3.5.1 Subject**

We analyzed data from a single rhesus monkey (monkey G, 7.1 kg, female) which was one of two animals studied in an investigation of MPTP (1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine)-induced parkinsonism. Here, we selectively analyzed data from the post-MPTP state. A prior report based on the pre-MPTP data [58] provides detailed explanations of several components of the experimental protocol that are also applicable to the post-MPTP dataset; for these components, we briefly recap the main points below.

### **3.5.2 Task**

Single-unit extracellular signals were collected during performance of a visually-cued reaching task, on which monkey G had been trained to a high degree of proficiency. Food rewards were delivered for successfully performed trials. Full task details are available in [58]. Here, the most relevant task component is the pre-reach delay period that spanned a variable-duration interval (1.01 - 3.50 s for the trials included in our analyses) from trial start to go-cue onset. During this delay period, the left hand rested at a start position (a metal bar) fixed at waist height to the

animal's left. The onsets/offsets of hand contact with this start position were detected with an infrared sensor.

### **3.5.3 Surgery**

Devices were implanted surgically to allow access to the right thalamus and globus pallidus via a parasagittal approach. The surgery was performed under sterile conditions with ketamine induction followed by Isoflurane anesthesia. Vital signs (i.e., pulse rate, blood pressure, respiration, end-tidal pCO<sub>2</sub>, and EKG) were monitored throughout the surgery to ensure a proper level of anesthesia. Analgesics and antibiotics were administered prophylactically.

On protocols.io, we have posted two step-by-step protocols relevant to the surgical procedures (craniotomy <https://www.protocols.io/private/8791490F352811EEAFC70A58A9FEAC02> and head fixation post and recording chamber implantation:

<https://www.protocols.io/private/A317A58E352711EEAFC70A58A9FEAC02>).

### **3.5.4 MPTP Administration**

The animal was rendered hemiparkinsonian by injection of MPTP into the right internal carotid artery (0.5-mg/kg, [59]). This model of parkinsonism was chosen to facilitate care of the animal and to increase the likelihood that the animal would perform the behavioral task following intoxication [60]. The MPTP administration procedure was performed under general anesthesia (1-3% Isoflurane) and prophylactic antibiotics and analgesics were administered post-surgically.

The animal developed stable signs of parkinsonism contralateral to the infusion (i.e., on the left side of the body).

On protocols.io, we have posted a protocol that includes a description of the procedures for internal carotid artery administration of MPTP

(<https://www.protocols.io/private/7BED34B9490A11EEACE60A58A9FEAC02>).

### **3.5.5 Data Acquisition and Spike Sorting**

Regions of GPi and VLa to be targeted for single-unit recording were located using standard neurophysiologic mapping methods [58]. The extracellular spiking activity of neurons in GPi and VLa was collected using multiple glass-insulated tungsten microelectrodes (0.5–1.5M $\Omega$ , Alpha Omega Co.) or 16-contact linear probes (0.5–1.0M $\Omega$ , V-probe, Plexon Inc.). Data were amplified (4 $\times$ , 2Hz–7.5kHz), digitized at 24 kHz (16-bit resolution; Tucker Davis Technologies), and saved to disk as continuous data. The stored neuronal data were high-pass filtered (Fpass: 300Hz, Matlab *firpm*) and thresholded, and candidate action potentials were sorted into clusters in principal components space (Off-line Sorter, Plexon Inc.; RRID:SCR\_000012). Clusters were accepted as well-isolated single-units only if the unit's action potentials were of a consistent shape and could be separated reliably from the waveforms of other neurons as well as from background noise throughout the period of recording. Times of spike occurrence were saved at millisecond accuracy.

On protocols.io, we have posted protocols relevant to electrophysiological recording (<https://www.protocols.io/private/7CF970AD3BE111EE882A0A58A9FEAC02>) and recording chamber maintenance

(<https://www.protocols.io/private/F56F59C63BDD11EE882A0A58A9FEAC02>).

### 3.5.6 Task-Aligned Analysis Windows

To focus on intervals of the monkey G spike trains when oscillatory activity was expected to be high, we restricted our analyses to task windows spanning the first second of each trial. These windows sampled data from the pre-reach delay periods; such intervals of enforced hand stillness have been historically associated with elevated beta power [61]. The 1000 ms window was chosen instead of the PSD segment size (1024 ms) because a small fraction of the go-cue onsets occurred at post-start latencies in (1000, 1024] ms. This selection step jointly served a second purpose of reducing potential sources of substantial non-stationarity in the analyzed spike data (along with a firing rate screening step, described further below).

### 3.5.7 Unit and Trial Selection Criteria

The process of selecting units and trials to include in analysis proceeded in three stages. As a first step, we excluded units recorded from the GPi and VL<sub>a</sub> when these failed to meet basic quality checks (e.g., poor isolation during spike sorting, or highly anomalous spike waveform features) or exhibited outlier FRs (as measured over the full recording period) of  $> 106.48$  Hz (GPi) or  $> 39.50$  Hz (VL<sub>a</sub>). Units with FRs  $< 1$  Hz were also excluded.

A second, trial selection step was prompted by the observation of units for which the “hold” (i.e., pre-reach delay) period-windowed data exhibited visible inhomogeneity in firing rates over trials, often appearing as blocks of trials during which spiking was either absent or very sparse, preceded or followed by blocks for which spiking was more regular. As this pattern was suggestive of considerable non-stationarity remaining in the hold period-restricted data, we implemented a

procedure to detect and remove windows associated with trials with aberrant firing rates. This procedure itself entailed multiple steps:

1. Prepare an  $nW \times 1$  vector  $S$  indicating the number of spikes in each 1000 s window.
2. Remove those rows from  $S$  reporting  $< 2$  spikes.
3. Use a Gaussian Mixture Model (GMM) to cluster the remaining entries of  $S$ , using a forward search procedure to set the number of Gaussians ( $\kappa$ ) parameter:
  - a. For each  $\kappa$  in 1 to  $\text{length}(S)$ :
    - i. Fit  $\kappa$  Gaussians to the data in  $S$  (Matlab *fitgmdist*) using initial centers ( $\mu$ ), variance ( $\sigma^2$ ), and mixing proportion ( $p$ ) settings as determined by the Variance Partitioning method [62].
    - ii. Record the Bayesian Information Criterion (BIC) value for this  $\kappa$ . The BIC is a negative log-likelihood-based measure that is penalized by the magnitude of  $\kappa$ ; lower BIC values are preferable.
    - iii. Break out of the loop when either:
      1.  $\text{BIC}(\kappa) > \text{BIC}(\kappa-1)$
      2. A failure to converge is reported for iteration  $\kappa$ .
  - b. Upon exiting the loop, return the GMM with the number of Gaussians parameter set to  $\kappa-1$ .
4. Assign each row in  $S$  to one of the  $\kappa$  Gaussians (hard clustering with Matlab *cluster*, which assigns the cluster labels that maximize the probability of observing the data, given the GMM parameters).
5. Find the Gaussian  $G$  with the maximum number of assigned windows. Retain only those windows from  $S$  that were assigned to  $G$ .



6. From the windows assigned to  $G$ , remove those windows for which the number of spikes  $nSpk$  is an outlier, given the mean and standard deviation parameters estimated for  $G$ . A window is defined as an outlier when  $|nSpk - \mu_g| > 3 * \sigma_g$ .

For each unit, the output of this trial (i.e., window) selection step was assumed to represent a set of 1000 ms segments for which the firing rates were relatively homogeneous. As a final step, we discarded those units for which  $< 30$  windows survived the trial selection procedure, or for which the mean firing rate over the surviving windows was  $< 1$  Hz (fraction of units discarded for the GPi: 6/98; VLa: 20/98). For the retained units, the median rates of trial removal were .99% (GPi) and 6.62% (VLa).

### **3.5.8 Windowed Data: Uncorrected PSD**

Note that Welch's method processes time series segments independently (i.e., the final estimated PSDs are the average of the individual PSD estimates for the Hamming-windowed 1024 ms data segments). Prior to calling *pwelch*, we prepared the hold period data by demeaning and adding 24 ms of zero padding to each 1 s hold period vector, and then concatenating these 1 s vectors into a single time series. This vector was then submitted to identical PSD analysis procedures as were described for the synthetic data.

### **3.5.9 Windowed Data: Shuffling-corrected PSD**

Similar to the synthetic dataset, many of the recorded units exhibited low FRs, such that local shuffling was not a practical option. We extended the global shuffling option from [7] (which was

only defined for unbroken time series) to temporally separated data segments with the following algorithm (all indices assume indexing starts at 1):

1. Input:  $dMat$  (“delta matrix”), an  $nD \times T$  matrix, consisting of the original  $nD$  delta vectors of length  $T$  msec (here,  $T = 1000$ ;  $nD =$  number of trials).
2. Create a list,  $sISI$ , consisting of the concatenation of all ISIs encountered in  $dMat$ , sorted in descending order of size. (Note that  $sISI$  will effectively function like a max-oriented priority queue.)
3. Initialize an  $nD \times T$  matrix  $wMat$  (“working matrix”) with zero entries.
4. Initialize an  $nD \times 1$  vector  $timeLeft$  with entries of the value  $T-1$  (here, 999), corresponding to the size of the largest ISI that might be inserted into each row of  $wMat$ . (Note that 1 ms must be reserved for insertion of the first spike for the first ISI).
5. While  $sISI$  is not empty:
  - a.  $isi \leftarrow$  top (max duration) entry, removed from  $sISI$
  - b.  $openRows \leftarrow$  indices of rows in which  $timeLeft \geq isi$
  - c.  $targetRow \leftarrow$  randomly chosen row from  $openRows$
  - d. Insert  $isi$  into  $wMat[targetRow]$ :
    - i. If this is the first ISI entered into  $targetRow$ , place “spikes” (1 entries) at  $wMat[targetRow][1]$  and  $wMat[targetRow][1+isi]$
    - ii. Otherwise, place a spike at  $wMat[targetRow][i+isi]$ , where  $i$  is the index of the rightmost spike present in the row
  - e.  $timeLeft[targetRow] \leftarrow timeLeft[targetRow] - isi$
6. For each row in  $wMat$ :

- a. If  $> 2$  spikes are present in a row, randomly re-sequence the ISIs (as in the standard shuffling procedure)
- b. If  $timeLeft[row] > 0$ : Right shift the spikes (insert a leading series of 0s) by a random integer drawn from  $[0, timeLeft[row]]$

7.  $shufMat$  (“shuffled matrix”)  $\leftarrow wMat$

To summarize, the algorithm implements global ISI shuffling by (1) initializing an empty data matrix, and a pool of ISIs to be inserted into it, (2) randomly inserting ISIs into the new matrix, proceeding from the longest to shortest ISIs, and packing the rows in a left-aligned fashion, and (3) randomly re-sequencing and right-shifting the ISIs in each row, to eliminate the structure introduced by step (2).

Each matrix of shuffled spikes was submitted to PSD analysis through an identical manner as was described for the uncorrected PSDs, and the standard deviation of the original PSD by the mean of the shuffling-generated PSDs was then used to produce the corrected PSD.

### 3.5.10 Windowed Data: Residuals-corrected PSD

Recall that, in the case of an unbroken time series, the residuals output,  $r(t)$ , is not defined for the first  $\hat{n}_r$  ms. In the case of the windowed data, the initial  $\hat{n}_r$  ms must be dropped for the residuals output corresponding to each trial. Therefore, the delta vectors corresponding to each individual trial were zero-padded by  $24 + \hat{n}_r$  ms, to both account for the initial RP-dependent lag, and to match the expected 1024 ms segment size. All other aspects of the residuals method remain as described in the context of the unbroken synthetic data.

### 3.5.11 Reference Synthetic Data for Comparison with Bursting Units

We conducted follow-up examination of two VLa spike trains, which exhibited spike timing patterns and corrected PSDs that were distinct from those generally observed in our synthetic datasets. To better characterize the distinctive features of these spike trains, which included suspected burst-firing behavior, we generated additional sets of reference spike trains, which did not include a bursting component, but did roughly match the two VLa units with respect to other key parameters. All simulations used the same framework as described above. For the synthetic data matched to the first VLa unit,  $T = 118 \times 1024$  ms,  $f_{osc} = 13$  Hz, and  $p_{base} = 15$  Hz. For the simulations matched to the second VLa unit,  $T = 51 \times 1024$  ms,  $f_{osc} = 14$  Hz, and  $p_{base} = 21$  Hz. For both sets of simulations,  $m = 0.6$ ,  $n_r = 1$ , and  $k = 0$ , and  $n = 100$  synthetic spike trains were generated.

### 3.5.12 Population-Level Comparisons of Oscillation Detection Rates

We used three analyses to compare the counts of significant oscillations reported by the two methods, for specific frequency bands of interest. All three analyses were performed over the pooled GPi and VLa data. The first analysis compared the proportions of PSDs for which the residuals and shuffling methods reported at least one point of significant power within the 8-30 Hz alpha-beta band. Proportions were compared with McNemar's  $\chi^2$  test, with continuity correction (R *mcnemar.test*, R Core Team, 2021; RRID:SCR\_001905). The second analysis focused specifically on those PSDs for which at least one method reported an alpha-beta oscillation (as defined above), and used logistic regression (Matlab *glmfit*, logit link) to ask whether the

corresponding unit's mean FR inversely predicted residuals-specific identification of alpha-beta incidence. The dependent measure was a binary coding of the PSDs (1 = alpha-beta detected by residuals method alone; 0 = detection by either both methods or shuffling only). The third analysis used McNemar's  $\chi^2$  to compare the proportions of PSDs for which the residuals and shuffling methods reported at least two consecutive points of significant power in the (0,4] Hz delta range. This delta band measure was intended to serve as a substitute for more traditional forms of 1/f-like trend assessment (which involve curve fits that do not accommodate possible RP distortion [63]).

## **4.0 One-State Residuals Method: Results**

### **4.1 Publishing Note**

The text and figures in this Chapter have been adapted from a manuscript that has been posted to a preprint server [19] and submitted to a peer-reviewed journal.

### **4.2 Results Overview**

In the following, we first consider the shuffling and residuals methods in the context of synthetic spike trains. We will highlight two example cases that aid in illustrating the procedures and outputs of the two correction methods, and then present the results of a systematic comparison of the methods' performance over simulations that encompass diverse parameter settings. Subsequently, we present the shuffling- and residuals-corrected output for the empirical single-unit data collected from the parkinsonian monkey, and highlight cases in which the divergence of the two methods' output is particularly notable.

## 4.3 Synthetic Data Results

### 4.3.1 Shuffling Method: Examples of Detected and Non-Detected Oscillations

Stochastic spike trains were generated according to the framework described in [2, 7]; full details are provided in the Materials and Methods. To briefly recap, spike arrivals were sampled from 1 kHz rate functions,  $p_{\text{spk}}(t)$ , which could be varied with respect to the parameters that governed the exponentially rising recovery period (steepness  $k$ ; duration  $n_r$ ), the oscillation (frequency  $f_{\text{osc}}$ ; modulation strength  $m$ ), the baseline FR ( $p_{\text{base}}$ ), and the simulation duration ( $T$ ). We considered a variety of parameter settings, although with an initial priority on relatively low  $p_{\text{base}}$  values, given the challenges to oscillation detection that are posed by these cases: The resulting low spike rates create a disproportionate reliance on the ISI distribution-based information that shuffling removes, since the opportunity for the oscillation to shape higher-order structure (e.g., serial ISI correlations) is reduced.

Figure 2 presents two illustrative cases generated from these simulation procedures (spike train duration  $T = 60 \times 1024$  ms,  $f_{\text{osc}} = 9$  Hz; for all synthetic data,  $n_r = 9$  and  $k = 0.7$ , unless otherwise specified). The two cases are similar in that standard statistical assessment of the uncorrected spike trains fails to detect the underlying 9 Hz oscillatory component, but differ in the factors that likely contribute to this failure, and in the shuffling method’s ability to recover this spectral feature.

Figure 2A shows the uncorrected spectra. As compared to the second (Figure 2A *right*) example, the first spike train was generated with a relatively high base FR (41 Hz), but with a somewhat low amplitude modulation factor ( $m = 40\%$ ). As noted in the Introduction, high FRs

lead to more dramatic RP-related distortions of the associated spectra [7]. Given these combined limitations (low  $m$ , greater distortion), the 9 Hz peak fails to cross the upper bound of the 95% corrected CI. The second spike train was generated with a stronger modulation setting ( $m = 60\%$ ), and a lower base FR (11 Hz), which results in less dramatic RP distortion; however, these advantages are offset by the reduced oscillation signal-to-noise ratio (SNR) that is also a known consequence of low FRs [2].

To both of these spike trains, we applied the standard global shuffling procedure (Figure 2B), yielding estimates of the renewal-equivalent power spectra (Figure 2A, red lines) and the shuffling-corrected PSDs. For the high FR, low  $m$  example, the mean shuffled data PSD follows the RP distortion trend while avoiding the implanted oscillation. Consequently, division by this shuffling-generated PSD achieves its intended outcome: The corrected spectrum (Figure 2C *left*) is generally flat, apart from the 9 Hz peak, which does surpass the upper frequency-derived CI threshold.

However, such successful peak recovery is not observed for the low FR, high  $m$  example (Figure 2C *right*). In this case, the shuffled PSD appears to capture both the distortion trend, and, to a modest extent, the 9 Hz peak. Correction reduces the relative height of this spectral feature, thereby compounding the challenges to detecting the spectral peak that were already present due to the low SNR. Even with the RP distortion removed, the 9 Hz peak falls short of the significance threshold.

Figure 2D, which plots the ISI PDFs for the two cases, provides a time domain perspective on these different shuffling outcomes, by highlighting the spike timing information that shuffling correction removes. In the higher FR, lower  $m$  case (Figure 2D *left*),  $> 50\%$  of the ISIs were  $\leq 23$  ms, thereby limiting the opportunity for the somewhat modest amplitude, 9 Hz oscillatory trend



(period  $\approx 111.11$  ms) to substantially influence the shape of the ISI distribution. Qualitatively, the PDF resembles the form expected for a Poisson process modulated by a recovery period alone (i.e., an initial dip followed by an exponential distribution). For the lower FR, higher  $m$  case (Fig 2D *right*),  $> 50\%$  of the ISIs exceeded 81 ms ( $\approx 73.8\%$  of the 9 Hz cycle length), and  $> 35.54\%$  were 112 ms or longer (i.e., the spiking skipped entire cycles). These long survival times afforded greater opportunity for the moderately strong oscillation term to more visibly shape the ISI distribution, with relative peaks and valleys appearing in rough correspondence to that expected for the 9 Hz rhythm.

To summarize, the lower FR/higher  $m$  example illustrates an instance relevant to the general caveat raised by [7]: If much of the rhythmicity of interest is present in the ISI distribution, then caution is warranted when applying the shuffling method (or the analytical alternative), as such correction will remove any oscillatory information present in the first-order ISI statistics. What our example here helps to clarify is that such scenarios emerge not only from precise pacemaker-like patterns, but also in spike trains in which a known sinusoidal drive is expressed through relatively sparse spiking. This observation motivated our interest in examining how the residuals method, which models and removes spike history effects restricted to a short history window, might perform in such scenarios.

#### **4.3.2 Residuals Method: Procedure and Example Case Performance**

Figure 3 illustrates the steps of the residuals method, and the outcomes for the two example synthetic spike trains introduced above. To reiterate, the main steps include (1) estimation of the RP duration, (2) fitting the spike train with the bounded last-spike point process model (PPM),

with the bound set to the RP duration, and (3) applying spectral analysis to the residuals of the fitted PPM.

#### 4.3.2.1 RP Estimation Algorithm

The RP duration estimation method is illustrated in Figure 3A-B (for the two example cases) and also Figure A5. (for a spike train with no oscillation). To recap (see Chapter 3 for full details), we sought to translate the general strategy outlined by [53] into a specific algorithm that might perform with reasonable effectiveness on limited, sparse spiking data.

We specifically implemented a procedure that compares the variance explained by a series of right-shifted exponential curves fit to the ISI PDF. We pursued this approach after observing, on cases such as the Figure 2-Figure 3 examples, that the curve goodness-of-fit often reaches its first local maximum when the exponential is anchored to the first post-RP time bin (i.e., when a unit with  $m = 0\%$  should first exhibit ISIs consistent with Poisson spiking), even when noise and oscillatory trends create visible departures of the post-RP region from a clean exponential form. In Figure 3A we highlight select iterations of the right-shifted curve fits, for the two examples. When the curves are left-anchored to the first possible start lag (1 ms), one observes visibly poor fits over the initial, RP-spanning time bins. As shown in Figure 3B, these mismatches correspond to relatively low values for our goodness of fit metric (the deviance difference, see Chapter 3). Iterating through the subsequent RP-containing bins, the fit measure gradually improves, reaching local peaks (in both examples) at the first post-RP time bin of 10 ms. After detecting subsequent decreases in the exponential term's explanatory power at 11 ms, we confirm the presence of initial local maxima, and return, for both spike trains, RP duration estimates ( $\hat{n}_r$ ) that match the ground truth  $n_r$  of 9 ms.

Note that more thorough evaluation of RP estimation performance is provided in the Methods Evaluation section, where we report on the systematic comparison of the shuffling and residuals methods.

### 4.3.2.2 Bounded Last-Spike Point Process Model

The next step of the residuals method entails the estimation of the shape of the RP (i.e., the  $t - s_*(t)$  effect; see Chapter 3) bounded within a historical interval of  $\hat{n}_r$  milliseconds. Fig. 3C illustrates the inputs and outputs of our current estimation procedure on the two example cases and with  $\hat{n}_r = 9$  ms. As was the case for global shuffling, the input (leftmost panels, zoomed in on the first 209 ms) was the entire binary spike train vector.

As stated in Chapter 3, we chose to use the point process modeling (PPM) framework to estimate the  $t - s_*(t)$  effect. In general, PPMs may operate on binary incidence vectors to estimate discrete-time approximations of conditional intensity functions (CIFs; see [6, 8] for additional background). When PPMs are used to predict future spikes on the basis of spike history, the CIF ( $\lambda[t|\theta, H(t)]$ ) represents the time varying probability that a spike will occur in time bin  $t$ , given the spiking history  $H(t)$  and the model parameters ( $\theta$ ). PPMs can be implemented in a variety of frameworks; we chose the well-established route of utilizing the Poisson Generalized Linear Model framework (i.e., linear regression with a log link function). The GLM took on the form  $\log\lambda(t|H_t, \beta) = \beta_0 + \sum_{j=1}^{\hat{n}_r} \beta_j \delta_{t-j, s_*(t)}$ , where  $\beta_0$  is the intercept term, the  $\delta_{t-j, s_*(t)}$  terms indicate whether the most recent spike fell within the time bin corresponding to  $j$  ms in the past, and the  $\beta_j$  weight the additive contributions of these last-spike lags to the predicted  $\log\lambda(t)$ . Once the vector  $\hat{\beta}$  of coefficient estimates is found, the linear expression is exponentiated to obtain the fit of the

model to the original spike train (excluding the first  $\hat{n}_r$  time bins for which no fit is defined). Due to the exponentiation step, the constant and the last-spike factors combine multiplicatively in the estimation of  $\lambda(t)$  [8].

Qualitatively, the GLM fits to the two example spike trains (Figure 3C, center panels) do appear to approximate the forms anticipated on the basis of the ground truth RP and  $p_{base}$  parameters. When no spikes occurred in the prior  $\hat{n}_r$  ms, the predicted  $\lambda(t)$  were approximately 0.041 and 0.011 for examples #1 and #2, respectively. When spikes did occur, the predicted  $\lambda(t)$  approximated the gradually rising RP trend.

#### 4.3.2.3 Spectral Analysis

Given a GLM fit, the residuals time series is generated as the actual-predicted time series (Figure 3C, rightmost panels; note the initial  $\hat{n}_r$  time bins are filled with zero padding). The residuals are submitted to identical spectral analysis steps as are applied to the original time series. For the two example cases (Figure 3D), the known 9 Hz peaks are successfully recovered, including for the lower FR spike train (Figure 3D, *right*) for which the shuffling method did not detect significant oscillation. This sensitivity gain – in tandem with the desired flattening of the RP distortion – is the result that we aimed to achieve by selectively removing bounded  $t - s_*(t)$  effects, and leaving intact the temporal structure expressed by the longer lags of the ISI distribution.

### 4.3.3 Methods Evaluation

#### 4.3.3.1 Oscillation Detection Performance: Low-to-Moderate Spike Rates

So far, we have shown a pair of example spike trains, for which the residuals method either matched or exceeded the shuffling method's detection of a known oscillation. To compare the two methods comprehensively, we ran a series of simulations and analyses that took into account a broad set of factors that may influence method performance, including the simulation parameter settings, and the corrected alpha level ( $\alpha_c$ ) of the significance test. Evaluation began with a primary synthetic dataset that focused on low-to-moderate base firing rates (range 8-64 Hz), oscillations in the theta-low gamma ranges (7-32 Hz), and the default recovery period parameters; we also generated a series of secondary datasets to address follow-up questions.

For the primary dataset, we directly manipulated four simulation parameters ( $T$ ,  $f_{osc}$ ,  $m$ , and  $p_{base\_offset} = p_{base} - f_{osc}$ ); the tables in Figure 4A-B list the full set of parameter settings. For each unique setting combination, we generated 100 spike trains, computed the shuffling- and residuals-corrected PSDs, and tabulated the true and false positive oscillation detection rates for each (see the Figure 4 caption for the hit rate (HR) and false alarm (FA) rate definitions).

To compare the classification accuracy of the two methods over a range of  $\alpha_c$  thresholds, we conducted a pROC (partial Receiver Operating Characteristic) analysis (Figure 4D). This analysis was implemented over a pooled dataset (i.e., collapsing over the parameter settings in Figure 4A-B, which are further considered in the additional analyses below). As in traditional ROC analysis, we plotted hit versus false alarm rates as the decision threshold ( $\alpha_c$ ) was varied, with shuffling and residuals curves generated for each of 1000 subsamples of the full spike train dataset (with the subsampling performed to obtain bootstrapped estimates of variability). Since the original curves

spanned variable FA and HR ranges, which did not fully span [0,1] (even with the inclusion of extreme significance thresholds), we truncated each curve to a partial ROC estimate that covered a restricted FA range (see the Figure 4 caption for details). The above steps enabled confirmation that the partial area under the ROC curve (pAUC) was reliably greater for the residuals method as compared against the shuffling method ( $t$ -test of  $pAUC_{res} - pAUC_{shuf}$ :  $t(999) = 1.288762e+03$ ,  $p \ll .001$ ). This result indicates that the residuals method generally improves oscillation detection performance on this collapsed-parameter, primary synthetic dataset.

We have described a rationale for anticipating that the residuals method may be especially advantageous for the lowest-FR spike trains, and in particular for sensitivity to the detection of oscillations when the base FR is low relative to the frequency of the oscillation to be detected. This reasoning leads to the prediction of an inverse relationship between  $p_{base\_offset}$  and the magnitude of the residuals-associated improvement in hit rates. A qualitative trend consistent with this prediction is visible in the top right table of Figure 4A (representing  $D_{HR}$  = the residuals-shuffling differences in HR, for the original, non-sampled data). For the select simulation parameters reported in the table, the predicted inverse relationship between  $D_{HR}$  and  $p_{base\_offset}$  is apparent as a general tendency for higher  $D_{HR}$  values (and warmer heatmap colors) towards the lower rows (corresponding to lower  $p_{base\_offset}$  settings) of the table. The Figure 4A  $D_{HR}$  values also point to a general inverted U-shaped trend relating  $m$  to  $D_{HR}$ , such that the residuals' advantage appears to be strongest for moderate modulation strengths, as compared to the lowest  $m$  levels (when the

oscillation is either non-existent or too weak for either method to detect) and highest  $m$  levels (when the oscillation is readily detected by both methods).<sup>6</sup>

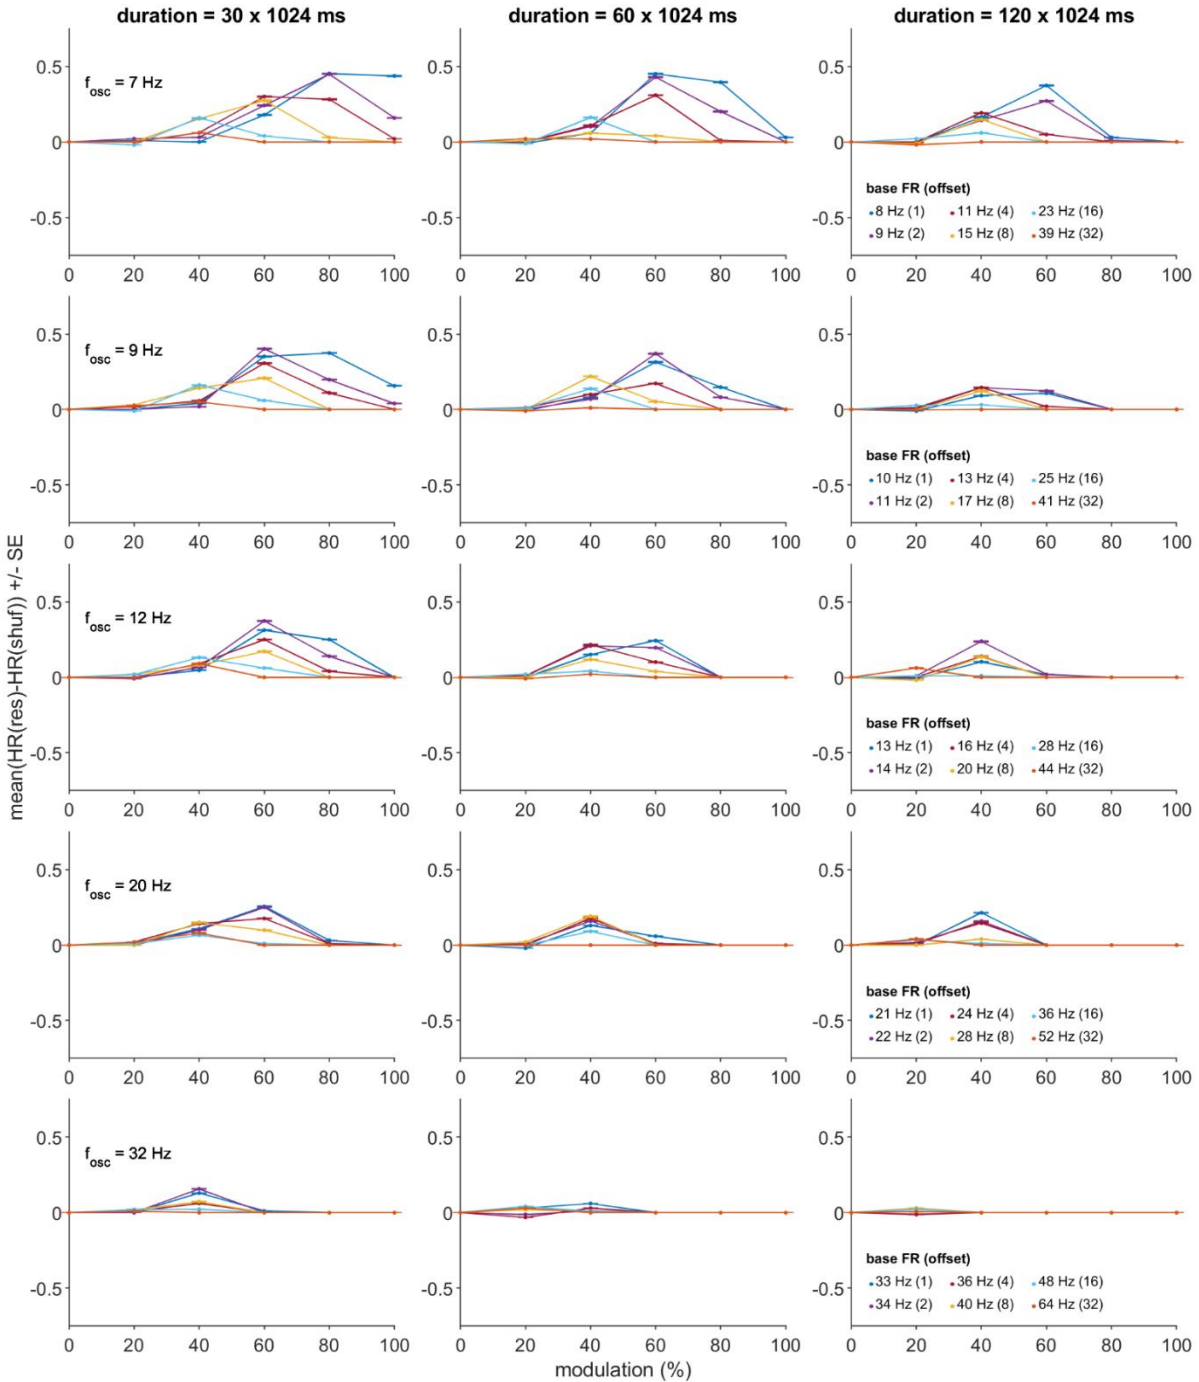
To quantify and statistically test these trends, and to assess the parametric effects on the methods' relative sensitivity and specificity more broadly, we conducted separate regression analyses of  $D_{HR}$  and  $D_{FA}$  ( $FA_{res}-FA_{shuf}$ ), with the significance threshold held fixed ( $\alpha_c = .05$ ), and the manipulated simulation parameters taken into consideration. As the regression analyses assume normally distributed noise in  $D_{HR}$  and  $D_{FA}$ , we conducted these regressions using  $D_{HR}$  and  $D_{FA}$  values derived from the same subsampled HR and FA output as informed the pROC analysis. Figure 5 and Figure 6 show  $D_{HR}$  and  $D_{FA}$  (summarized as the mean  $\pm$  standard error over subsamples) for all 540 unique combinations of the manipulated parameters. The  $D_{HR}$ - and  $D_{FA}$ -specific regression models included predictors corresponding to all four directly manipulated parameters ( $T$ ,  $f_{osc}$ ,  $p_{base\_offset}$ , and  $m$ , with all  $m = 0$  cases excluded), and a quadratic term for  $m$ .

The full regression results for the primary synthetic dataset are provided in Tables A1-A2; we highlight a subset of the significant effects here. The  $D_{HR}$  results (Table A1) statistically confirmed the predicted inverse relationship between  $D_{HR}$  and  $p_{base\_offset}$  ( $p_{base\_offset}$  term:  $t = -292.25$ ,  $p \ll .001$ ) and the observed inverted U-shaped relationship between  $D_{HR}$  and  $m$  ( $m^2$  term:  $t = -296.41$ ,

---

<sup>6</sup> We can make an additional notable observation regarding the 30 x 1024 ms (= 30.72 s) column in Fig. 4B: When presented with beta band frequencies (12 Hz, 20 Hz) and base firing rates that modestly exceeded these frequencies (14 Hz, 22 Hz), the shuffling method failed to detect the oscillation on only 33% (20 Hz) or 62% (12 Hz) of trials. This observation may be helpful to keep in mind when interpreting existing studies that use the shuffling method, since experimental constraints can at times require researchers to rely on such short samples of spike data (e.g., see the intraoperative data reported in Chan et al. (2010) [65]).

synthetic data: hit rate differences (residuals - shuffling), low to moderate FRs (9 ms RP,  $k = 0.7$ )



**Figure 5. Residuals (res) - shuffling (shuf) difference in hit rates ( $D_{HR}$ ) over the varied parameters of the primary synthetic dataset.** Means and standard errors (SE) reflect summaries over the hit rates (HR) computed for each of 1000 subsamples of the original, primary synthetic dataset, which was generated using low-to-moderate firing rates (FR) and a 9 ms relative recovery period (RP;  $k =$  steepness parameter). See Chapter 3 and the Figure 4 caption for details regarding the definition of a hit and the

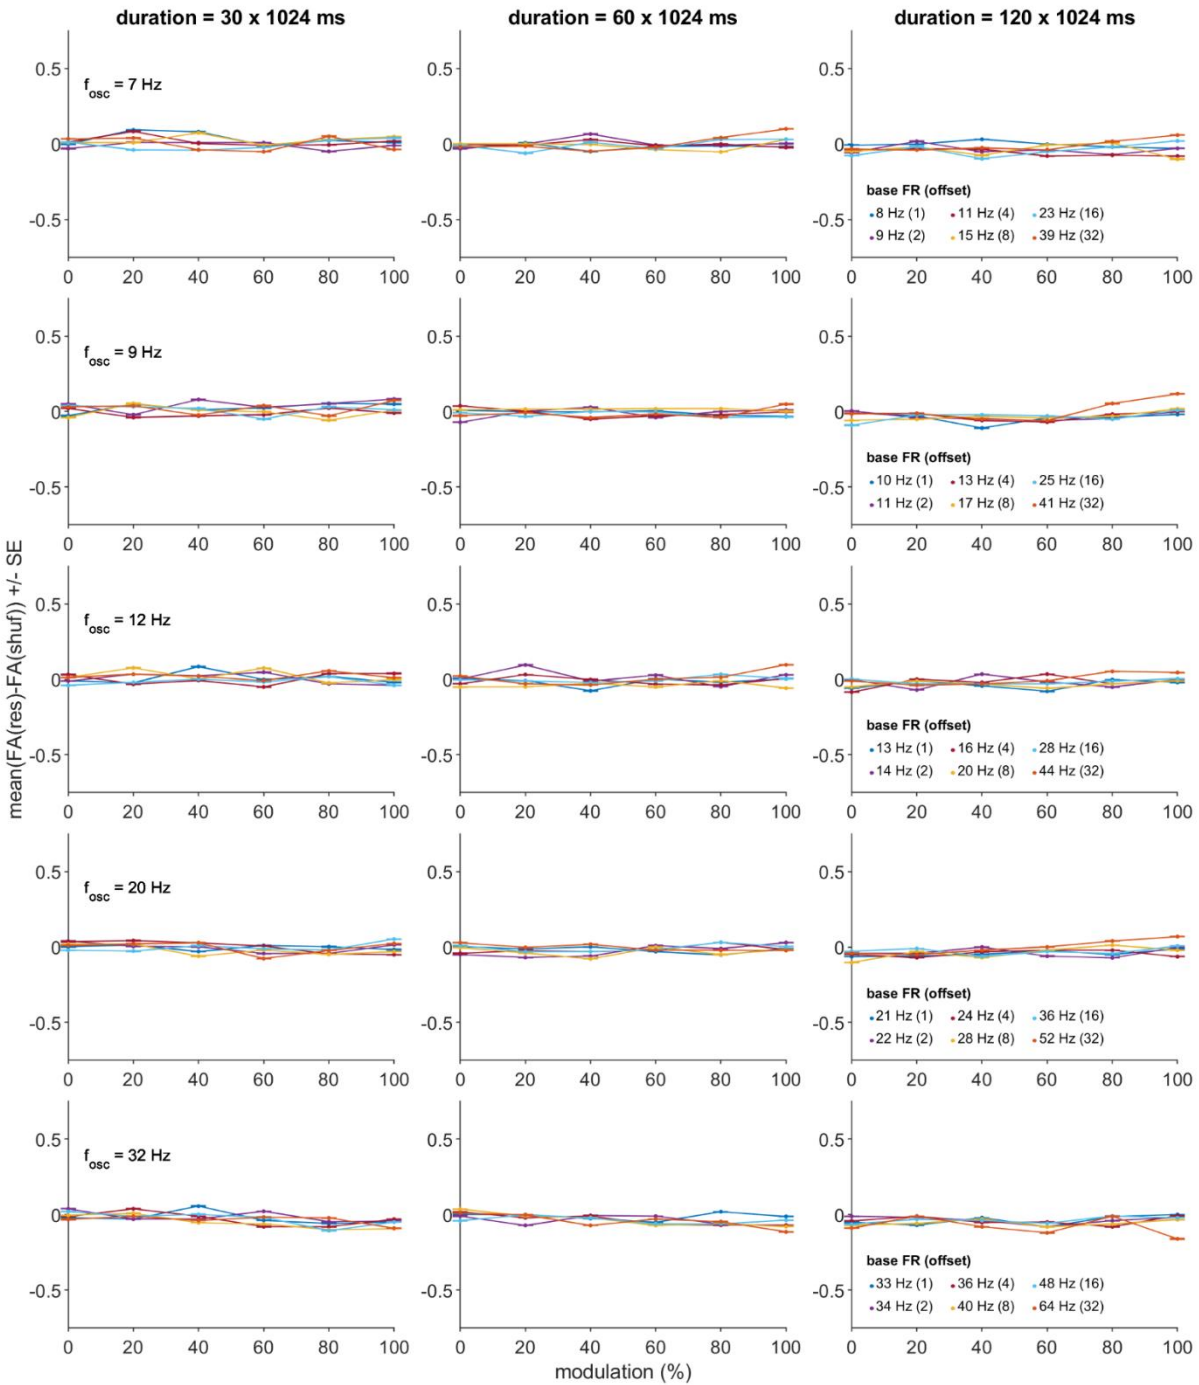


subsampling procedure. Plots depict all 540 unique combinations of oscillation frequency ( $f_{osc}$ , rows), simulation duration ( $T$ , columns), oscillation modulation strength ( $m$ ,  $x$  axes) and base FR - oscillation frequency offset ( $p_{base\_offset}$ , lines).

$p \ll .001$ ). As an additional test of the  $p_{base\_offset}$  effect – which was correlated with  $p_{base}$  in these simulations, thereby creating ambiguity in the interpretation of the effect as reflecting either base FR -  $f_{osc}$  or the absolute base FR – we created a follow-up synthetic dataset in which a directly manipulated  $p_{base}$  factor (levels = 15 or 20 Hz) was crossed with an independently varied  $p_{base\_offset}$  factor (levels = [2:2:8] Hz) and the same  $T$  and  $m$  parameter sets as were described above. Regression analyses of the  $D_{HR}$  estimates from these simulations (Table A3;  $D_{HR}$  values are plotted in **Error! Reference source not found.**) replicated the expected negative  $p_{base\_offset}$  effect ( $t = -59.81$ ,  $p \ll .001$ ). This result supports the interpretation of the margin between the firing rate and to-be-detected oscillation frequency as independently influential on the residuals method's enhanced sensitivity (beyond the effect of the absolute  $p_{base}$ ).

No specific predictions were made regarding the two methods' relative false alarm rates, and for many of the parameter combinations represented in Figure 6 (and **Error! Reference source not found.** for the dataset with the direct  $p_{base}$  manipulation), the  $D_{FA}$  values clustered near zero. Nevertheless, the very high  $n$  of subsamples considered resulted in high sensitivity to parametric effects, and significant main effects and interactions were found (Table A2; Table A4); however, their modest scale leaves uncertain the practical implications of these trends. As discussed below, more noticeable trends were found upon introducing higher FRs to the simulations.

synthetic data: false alarm differences (residuals - shuffling), low to moderate FRs (9 ms RP,  $k = 0.7$ )



**Figure 6. Residuals - shuffling difference in false alarm rates ( $D_{FA}$ ) over the varied parameters of the primary synthetic dataset.** Means and standard errors reflect summaries over the false alarm rates (FA) computed for each of 1000 subsamples of the original dataset. See the Figure 4 caption for details regarding the definition of a false alarm. Abbreviations, plotting conventions, and the subsampling procedure follow from those described for Figure 5 and Figure 4.

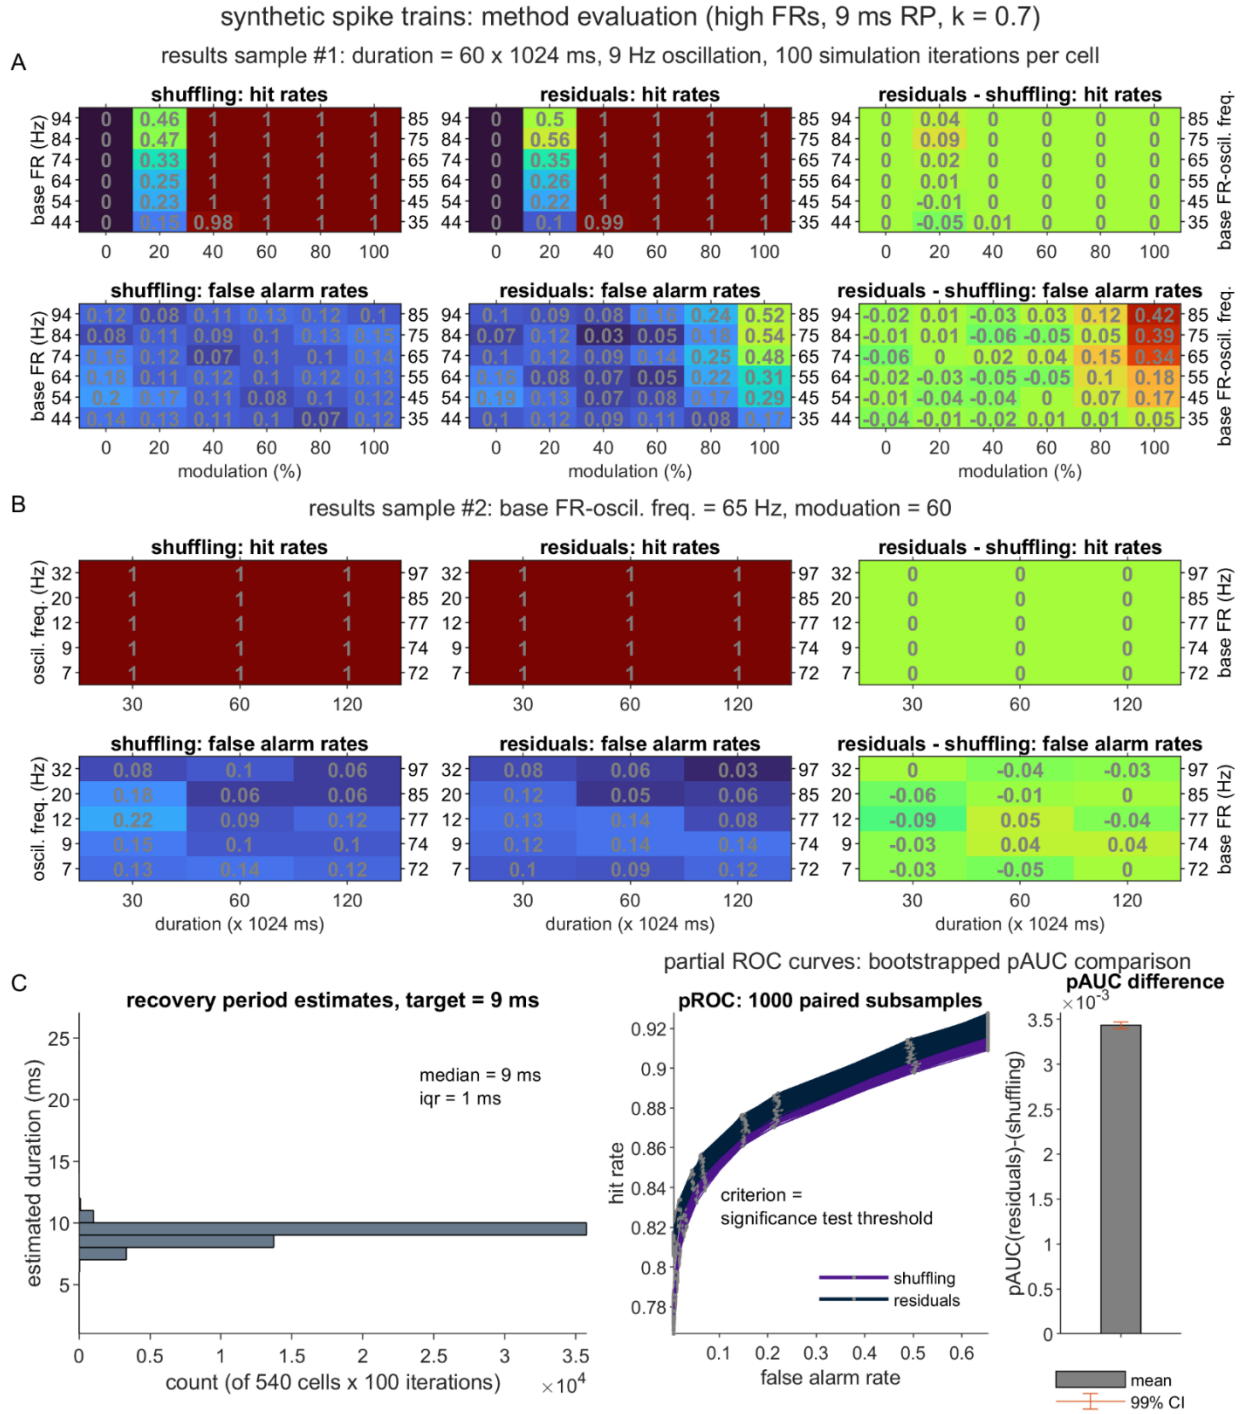
#### 4.3.3.2 Oscillation Detection Performance: High Firing Rates

The residuals method was motivated by a problem presented by sparse spike trains, but we might also ask how effectively it performs with more vigorously spiking units. To this end, we created a secondary dataset that differed from the primary dataset in its use of higher  $p_{base\_offset}$  values (Figure 7-9).

Three patterns stand out in the results. First, as anticipated, high FRs reduced the residuals method's overall performance advantage:  $pAUC_{res} - pAUC_{shuf}$  remained reliably positive (Figure 7D), but with only a narrow margin of performance improvement. Second, as also anticipated, this narrowed performance gap was mirrored by a similar narrowing in the residuals' sensitivity advantage.  $D_{HR}$  values remained mostly positive but very modest (Figure 8; Table A5).

Third, we observed a complex interaction pattern in the methods' relative false alarm rates. Complete illustrations and statistical results are available in Figure 9 and Figure A8.-

Figure A11., and Table A6. Here, we will highlight four principal observations (all based on quantities summarized over the 1000 subsamples; see Chapter 3). (1) As duration  $T$  increased, and especially as the modulation strength  $m$  approached 100%,  $D_{FA}$  scores grew either very positive (for low oscillation frequencies) or negative (for high frequencies; Figure 9). In other words, in the high  $T$ , high  $m$  regime, the residuals method switched from producing more FAs to fewer FAs than shuffling did, as  $f_{osc}$  increased ( $T \times f_{osc} \times m$  interaction:  $t = -178.70$ ,  $p \ll .001$ ). (2) Although  $f_{osc}$  is



**Figure 7. Performance of the shuffling and residuals methods over a synthetic dataset of high firing rate (FR) spike trains.**

Panels (A)-(D): Plotting conventions, hit and false alarm definitions, and analysis procedures are identical to those described for the primary dataset depicted in Figure 4. Relative to the primary dataset, this high FR dataset differed in the use of greater  $p_{base\_offset}$  values (see “base FR - oscil. freq.” tick labels in Panel (A)).

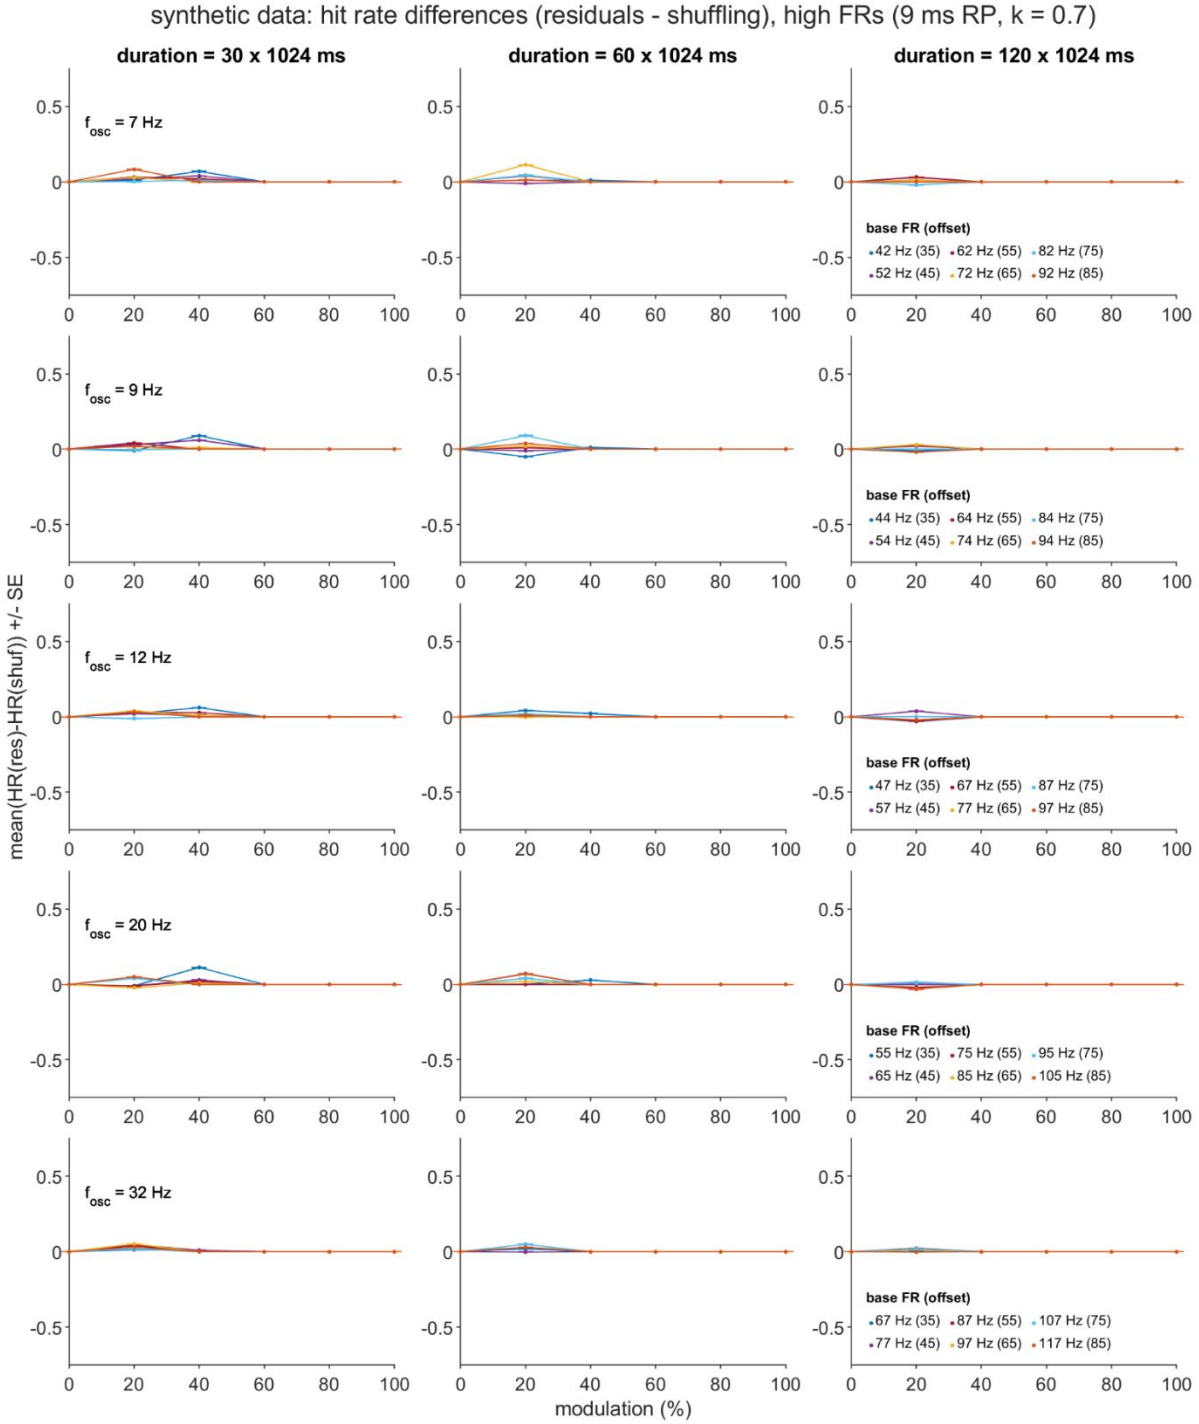
correlated with the  $p_{base}$  parameter, the sign flip in  $D_{FA}$  scores cannot be easily attributed to increasing FRs. In the  $T = 120 \times 1024$  ms,  $m = 100\%$  cases, the effect of  $p_{base\_offset}$  on  $D_{FA}$  negatively interacted with  $f_{osc}$  ( $t_{f\_osc \times p_{base\_offset}} = -68.23$ ;  $p \ll .001$ ), reflecting a tendency for the sign of the FR offset effect to likewise flip. (3) Plots of the average absolute FAs for the lowest  $f_{osc}$  setting (7 Hz, Figure A8.) show that in the high  $T$ , high  $m$  cases, the shuffling FA rates remained mostly low (mean over all subsample  $\times p_{base}$  combinations = 10.16%) in contrast to the elevated residuals FA rates (overall mean = 47.55%). For the highest  $f_{osc}$  setting (32 Hz, Figure A9.), the high  $T$ , high  $m$  cases resulted in very elevated FA rates for shuffling (overall mean = 76.09%) and less dramatically elevated FA rates for the residuals method (overall mean = 31.73%). (4) Summarizing over all  $T = 120 \times 1024$  ms,  $m = 1.0$  cases, both methods' false detections were typically characterized by very low peak power, relative to that observed for hit points (mean  $power_{FA}/power_{Hit} = 0.15$  for residuals and 0.22 for shuffling), and by high, gamma range frequencies (means = 83.10 Hz for residuals and 57.86 Hz for shuffling). Although an ideal solution would prevent the false detections, these consistent features could aid in the development of post-hoc rules for flagging suspected FA points in either method's corrected PSDs.

To illustrate the above observations, **Error! Reference source not found.**

**Figure A11.** present individual corrected PSD examples (in addition to samples of the associated synthetic spike trains) drawn from high FR, 100% modulation strength cases. **Error! Reference source not found.** compares example residuals- and shuffling- corrected PSDs for the  $f_{osc}$  levels that resulted in the most elevated false alarm rates for each method (7 and 32 Hz, respectively).

**Figure A11.** compares two residuals PSDs for these same two

$f_{osc}$  levels, with matched 100 Hz base FRs (a setting not included in the above high FR dataset,



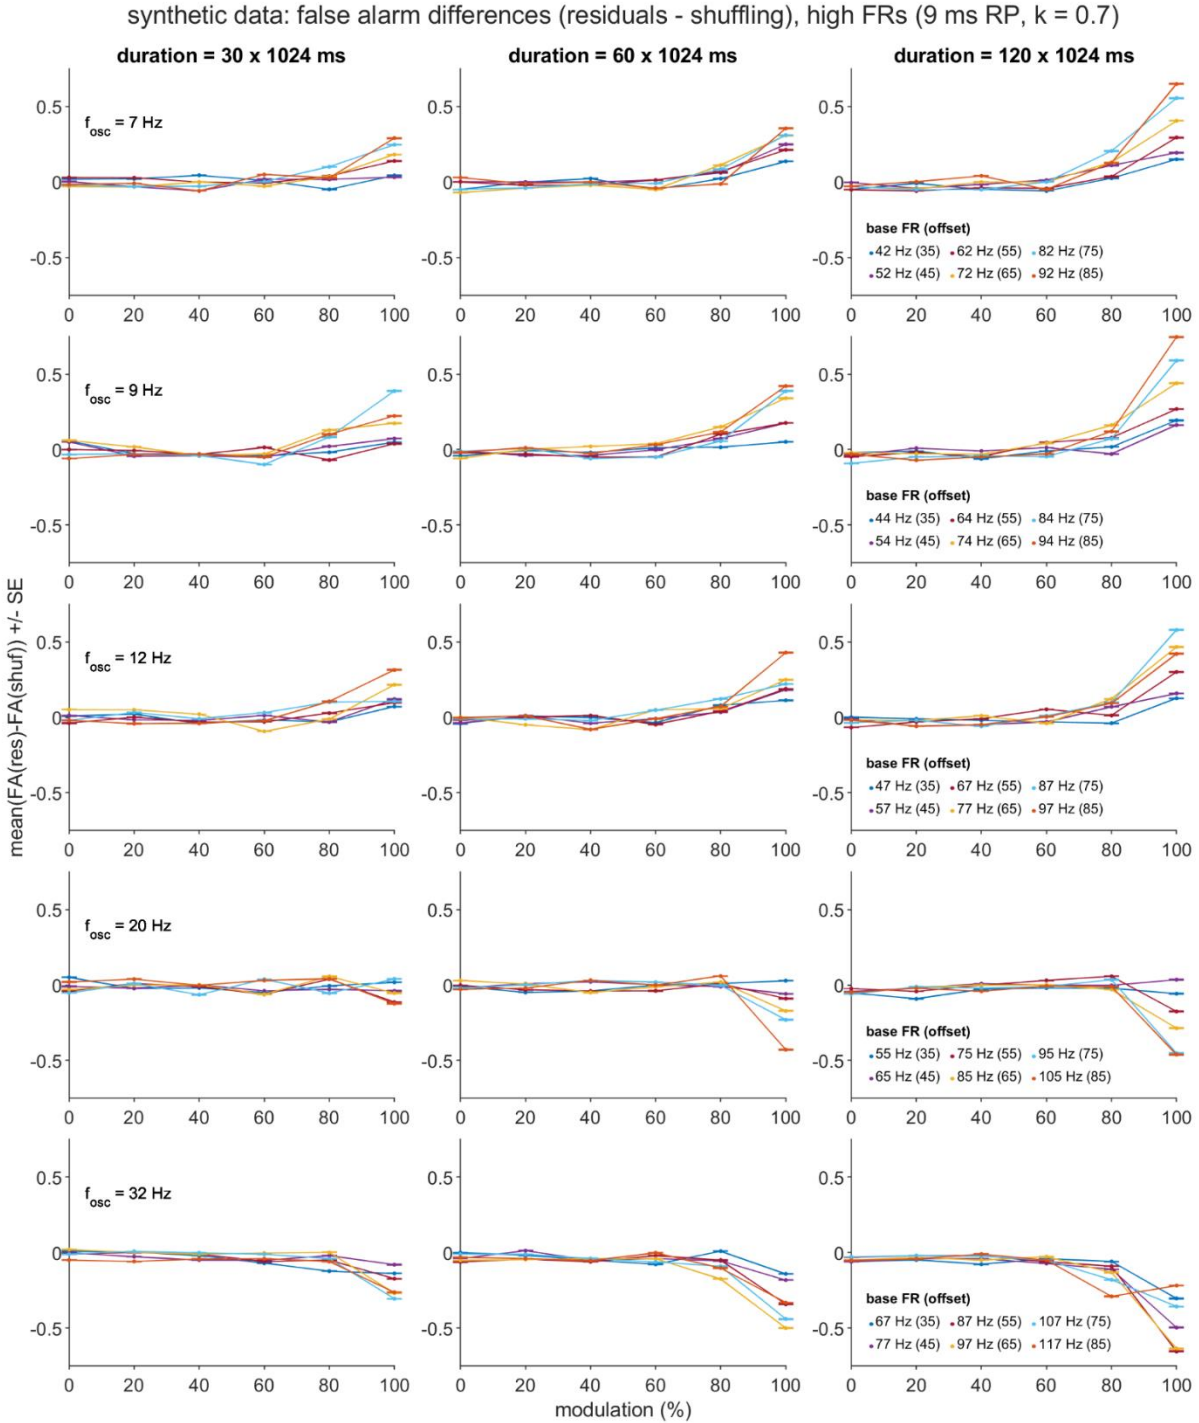
**Figure 8. Residuals - shuffling difference in hit rates over the varied parameters of the dataset of high firing rate spike trains.** Means and standard errors reflect summaries over the hit rates computed for each of 1000 subsamples of the original, high FR dataset (see **Figure 7** for details). Abbreviations and plotting conventions follow from those described for **Figure 5**.

but added here to illustrate the  $f_{osc}$  effect across matched base FR levels).

#### 4.3.3.3 Evaluation of Recovery Period Duration Estimation

Overall, the preceding analyses indicate that the residuals method achieves its primary objective of enhancing oscillation detection in sparse spiking scenarios, while maintaining performance comparable to that of the shuffling method when spiking is more vigorous (excluding the limited high-FA cases described above). Correspondingly, the first step of the method – the estimation of the RP duration – was also reasonably accurate. In the primary synthetic dataset, for the RP estimates collapsed across the full set of 540 x 100 spike trains (Figure 4C), 51.07% identified the 9 ms  $n_r$  setting precisely, 85.13% erred by  $\leq 1$  ms, and 94.78% erred by  $\leq 2$  ms. The follow-up, high FR dataset also demonstrated generally low error rates (Figure 7C: % of absolute errors,  $|err|$ ,  $\leq 0$ , 1, or 2 ms: 66.22%, 93.53%, 99.87%).

To determine whether this accurate performance might generalize to other RP durations and forms, we generated three additional secondary datasets, with the original RP parameters modified to produce either a long relative RP ( $n_r = 18$  ms, steepness parameter  $k = 0.7$ ), a short, sharply-rising relative RP ( $n_r = 4$  ms,  $k = 0.4$ ), or a short absolute RP ( $n_r = 3$  ms,  $k = 0$ ). Estimates clustered closely around the ground truth  $n_r$  for both the 3 ms absolute RP (**Error! Reference source not found.**12C: %  $|err| \leq 0$ , 1, or 2 ms: 87.28%, 96.62%, 98.90%) and the 4 ms relative RP (Figure A13.13: %  $|err| \leq 0$ , 1, or 2 ms: 84.03%, 96.58%, 99.01%). When the relative RP with the default  $k = 0.7$  steepness is doubled to 18 ms, absolute error rates are comparable to those for the original 9 ms duration if the deviations under consideration are doubled accordingly (Figure A14.4C: %  $|err| \leq 0$ , 2, or 4 ms: 44.79%, 86.73%, 92.72%).



**Figure 9. Residuals - shuffling difference in false alarms over the varied parameters of the dataset of high firing rate spike trains.** Means and standard errors reflect summaries over the false alarm rates computed for each of 1000 subsamples of the original, high FR dataset (see **Figure 7** for details). Abbreviations and plotting conventions follow from those described for **Figure 6**.



These datasets also afford the opportunity to ask whether the residuals method retains its performance advantage when confronted with these differing RP features<sup>7</sup>. As reported above for the default RP datasets, panels A, B, and D of **Error! Reference source not found.**-Figure A14. show, for each RP setting and correction method, highlighted HR and FA outcomes for select parameter combinations, and the overall pROC and pAUC output. For all three RP configurations, residuals remains the better-performing oscillation detector ( $t$ -tests of  $pAUC_{res}-pAUC_{shuf}$ :  $ts(999) > 1.0608e+03$ ,  $ps \ll .001$ ).

## 4.4 Experimental Data Results

### 4.4.1 Motivation and Approach

Further comparison of the residuals and shuffling methods, conducted on real, biological single-unit data, served two purposes. First, any qualitative discrepancies between the neuronal and synthetic spectra provide useful information that can help inform the development of more accurate simulation frameworks in the future. Second, and most importantly for the present purposes, these evaluations can help experimentalists anticipate what might be encountered when

---

<sup>7</sup> To clarify: These tests assessed the robustness of the residuals method's performance advantage when the RP settings were varied across different synthetic spike train datasets. I do not test the method's handling of hypothetical RP features that might vary over the course of a recording, since the current implementation of the method assumes a fixed form for the RP, given a single spike train instance.

applying either PSD correction method to their own data, and determine what further development, if any, might be needed to tailor the methods to the specific needs of real-world research questions.

In empirically-observed, neuronal spike trains, the ground truth is unknowable. That said, we prioritized sampling neural data acquired during circumstances for which the prior expectation of oscillations is high, based on evidence from sources that are distinct from spike trains. Towards that end, we evaluated the two correction methods' output on single-unit data recorded from the internal globus pallidus (GPi) and GPi-targeted ventrolateral anterior thalamus (VLa) of a monkey with MPTP-induced parkinsonism, during the pre-reach delay (i.e., "hold") periods of a cued reaching task. Across multiple species and motor-related brain regions, pathologically exaggerated oscillations in neural activity are a hallmark of parkinsonism, especially within the beta (12-30 Hz) and neighboring alpha (8-12 Hz) frequency bands (see [64, 65] for recent general reviews, and [66, 67] for examples specific to the thalamus and GPi, respectively). Moreover, alpha-beta oscillations are especially pronounced during periods of postural maintenance, such as the hold periods under consideration here [61]. Since the bulk of the evidence concerning these oscillations has been obtained from recordings of local field potentials (LFPs), our expectation of their occurrence in spike trains is based on a source that, while ultimately derived from the same underlying electrophysiological signal, is not affected by RP-related (and general point process-related) challenges to oscillation measurement.

The shuffling and residuals methods' output were compared for 92 GPi units (analyzed recording duration min, median, max = 30, 80.5, 262 s) and 78 VLa units (31, 74.5, 191 s). As anticipated based on previous reports [58], firing rates for the analyzed rest periods were generally slower for the VLa than for the GPi ( $p \ll .001$  by the Wilcoxon rank sum test; min, median, max FR for VLa = 3.09, 12.24, 41.02 Hz; for GPi = 3.13, 39.90, 114.96 Hz). Our simulation results

allow us to predict the likely sensitivity of the two methods for detecting oscillations in data with durations and firing rates typical of this empirical dataset. The synthetic data hit rate results for a representative low beta oscillatory frequency (12 Hz), low base firing rate (13 Hz), and short spike train duration (30 x 1024 ms; **Error! Reference source not found.**) predict that, for an oscillation modulation strength of at least 60%, the residuals method at minimum will detect the oscillation > 65% of the time.


Application of the two methods to the empirical dataset revealed considerable variability in the resulting corrected PSDs, but some recurring patterns do stand out. Here, we will highlight two of these patterns with four representative units (Figure 10, Figure 11). The full population results are summarized in Figure 12. To highlight the (0,100] Hz region over which we sought significant power, all neural spectral plots are confined to the initial 100 Hz; however, note that full spectra may be generated using code from the Github repository that accompanies this report.

#### 4.4.2 Single Unit Examples

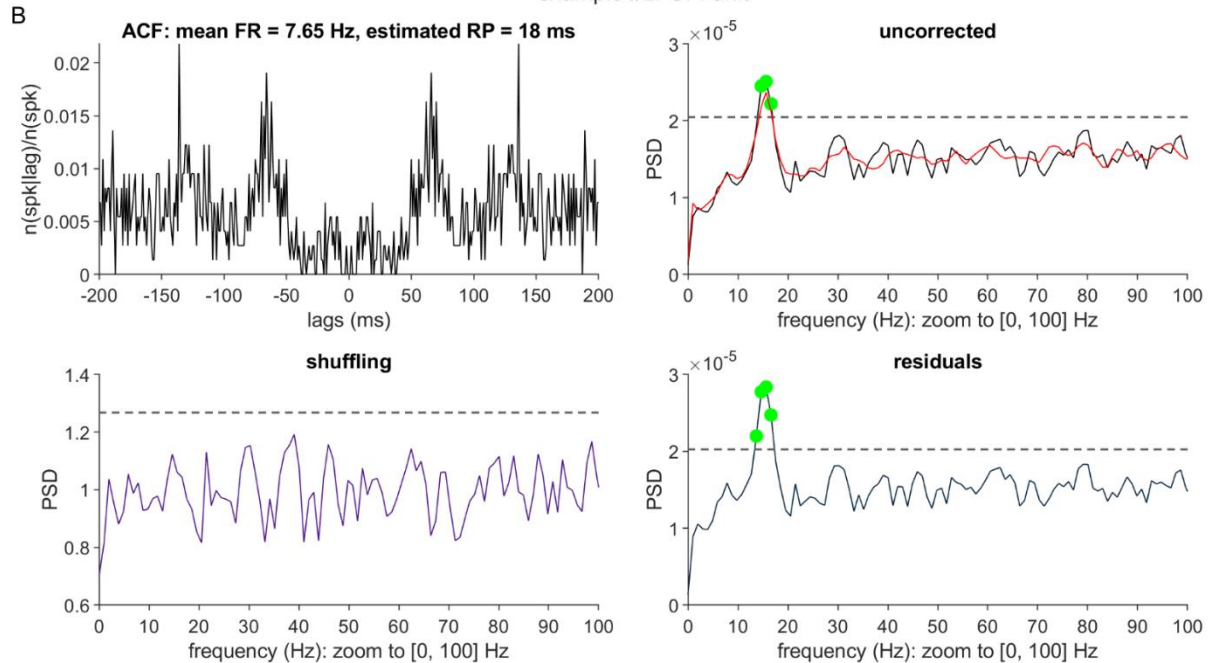
Figure 10 presents two units that illustrate the first recurring pattern, which is characterized by its consistency with what the simulations would predict: beta (or alpha) peaks that were recovered by the residuals method exclusively, from sparse spike trains that share qualitative features with our synthetic data. Figure 10A shows such an example from the VLa (mean FR = 11.90 Hz, analyzed duration = 63 s). The uncorrected PSD matches the canonical RP distortion form, with a candidate beta oscillation (peak = 13.67 Hz) sitting atop the depressed low-frequency region. The oscillation is flagged as significant in the residuals PSD, but not the uncorrected or shuffling PSD. In this case, the behavior of the shuffling method is similar to that illustrated in the

post-MPTP NHP spike trains: candidate  $\beta$  oscillations (axis bounds vary to aid visualization)

example #1: VL<sub>a</sub> unit



example #2: GPi unit



**Figure 10. Comparison of the shuffling and residuals output for two highlighted units acquired from a parkinsonian non-human primate (NHP).** Spike trains were sampled from one monkey that had been injected with the MPTP (1-methyl-4-phenyl-1,2,3,6-tetrahydropyridine) neurotoxin, and restricted to resting periods of a reaching task. Example units exhibit autocorrelation functions (ACF) consistent with a possible  $\beta$  (beta band; 12-30 Hz) oscillation. (A) Example unit from the ventrolateral anterior thalamus (VL<sub>a</sub>). Top, left: ACF, normalized by total spike count; spk = spike. Top right, bottom left, and bottom right: Uncorrected,

shuffling-corrected, and residuals-corrected PSDs for the VLa unit, following the formatting and statistical conventions from Figure 2-**Figure 3**. In the uncorrected spectrum panel, the ISI-attributed PSD (red line) was estimated via an adaptation of the global shuffling method to temporally separated task windows (see Chapter 3). (B) The same sequence of plots as described in (A), as applied to a unit from the internal globus pallidus (GPi).

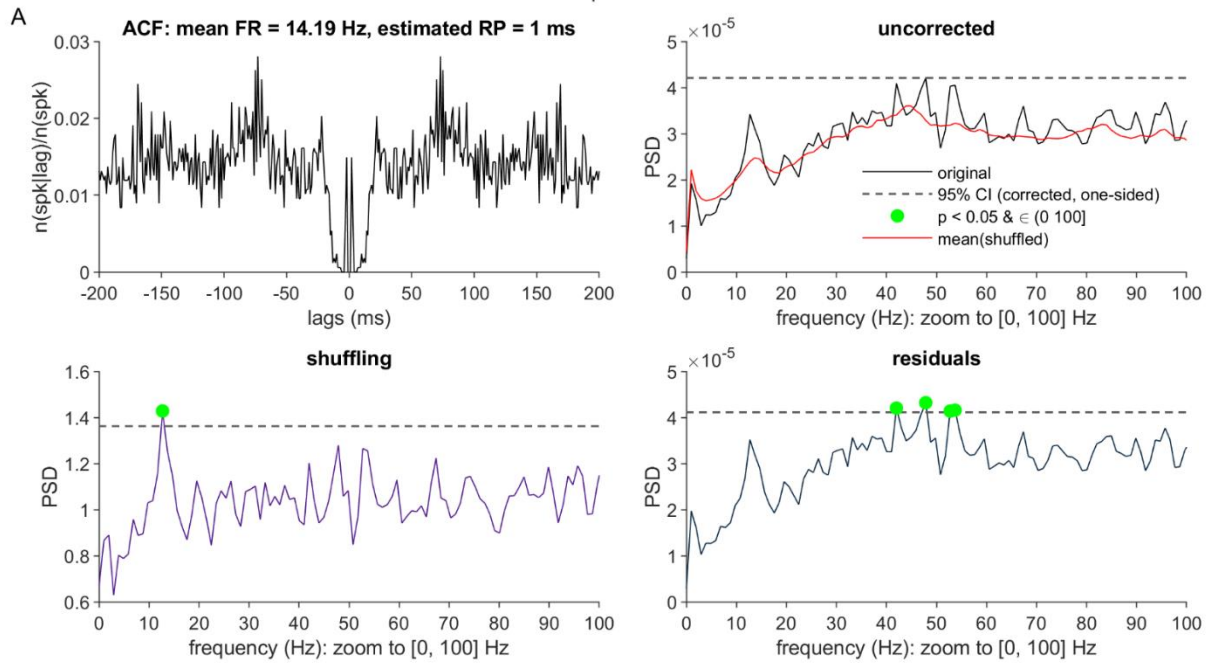
right-most column of Figure 2A: The shuffling approximation of the renewal equivalent PSD captures not only the global distortion-consistent trend, but also much of the candidate beta peak, resulting in partial removal of this feature from the shuffling-corrected PSD. The 11.90 Hz FR indicates that, when considered relative to the candidate beta oscillation's frequency, a spike occurred on average just .87 times per cycle. Figure 10B shows a second example of this pattern, in an unusually slowly-spiking GPi unit (mean FR = 7.65 Hz, analyzed duration = 96 s, peak of alpha-beta oscillation = 15.62 Hz).

Figure 11 presents two VLa units that illustrate the second pattern, broadly categorized as suspected alpha-beta oscillations accompanied by temporal structure that was not incorporated into our simulations. Such additional structure included both burst firing (that is, transient intervals of elevated spike rates, see both panels A and B) and prominent 1/f-like trends (panel B). The two correction methods' behavior could differ considerably in these scenarios, with variable outcomes for their relative sensitivity to likely oscillations.

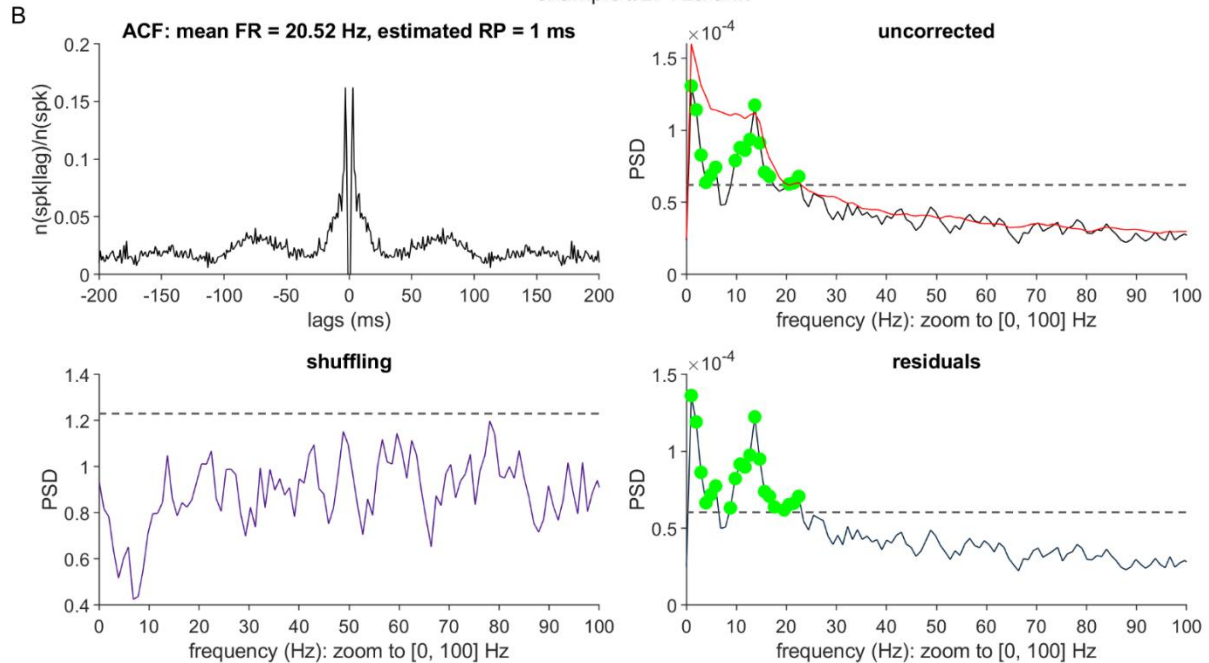
In the first VLa example (Figure 11A), a likely low beta oscillation (peak alpha-beta frequency = 12.70 Hz) appears in tandem with a sporadic tendency for the unit to fire in a highly stereotyped pattern of two-spike bursts. The joint presence of these two patterns is most apparent in the autocorrelation plot (top left panel). Key features of the intra- and extra-burst spiking can be discerned from the ISI PDF (Figure A16.A) and joint  $[ISI_n, ISI_{n+1}]$  PDF (Figure A16.C, [68]): ISIs of 2-3 ms occur exclusively between pairs of spikes that are bracketed by much longer ISIs ( $\geq 14$

post-MPTP NHP spike trains: candidate  $\beta$  oscillations (axis bounds vary to aid visualization)

example #1: VLa unit



example #2: VLa unit



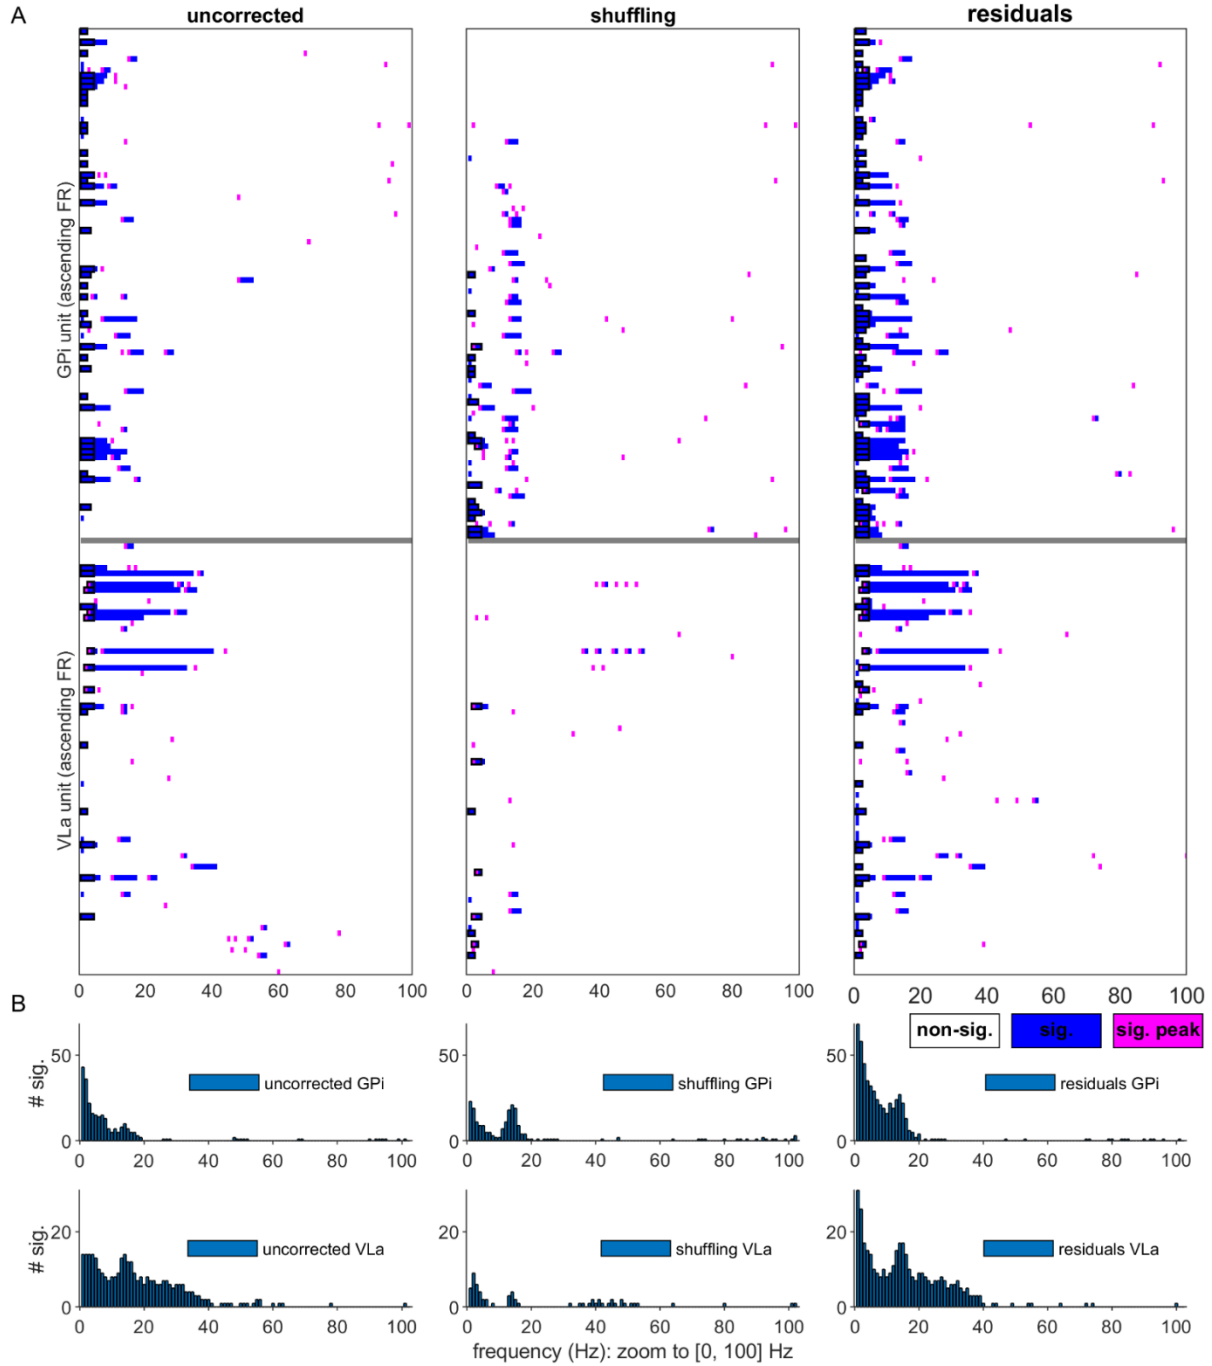
**Figure 11. Comparison of the shuffling and residuals output for two units with putative beta oscillations and non-oscillatory features.** Spike trains originated from two ventrolateral anterior thalamus (VLa) units, which had been recorded from the same parkinsonian non-human primate (NHP) that contributed to Figure 10 and Figure 12. Panels (A)-(B): Abbreviations and the statistical and plotting conventions follow from those described for Figure 10.

ms), and all remaining spikes are separated by ISIs that last at least 7 ms (and often much longer). As such, the ISI statistics for this unit are consistent with a description of the spike train as switching between two different states, for which the effective recovery period varies from very short (intra-burst) to much longer (extra-burst).

This two-state characterization lends insight into the PSD outcomes. In the uncorrected spectrum, the beta peak is sunken into a typical distortion-related trough, and the shuffling method achieves satisfactory flattening of the spectrum and raising of the peak. In contrast, the residuals method performs very little distortion correction, due to the 1 ms estimated RP duration (see Figure A16.6A-B, *left*) and resulting 1 ms bound on the last-spike effects estimated by the PPM. Such a model should adequately account for RP effects within the two-spike bursts: On simulated, non-bursty spike trains with a 1 ms recovery period and parameters that otherwise roughly matched those for this VLa unit (see Chapter 3), the residuals method exhibited accurate oscillation detection (for  $f_{osc} = 12$  Hz;  $m = 0.6$ : hit rate = 100%, FA rate = 7%). The apparent under-correction in this VLa example is likely due to the insufficient removal of the long RP effects that occur outside of the bursts. In Chapter 5 I propose an extension upon the residuals method that may address the need to accommodate burst and non-burst states with differing RP dynamics.

In the second VLa example (Figure 11), a likely low beta oscillation (peak alpha-beta frequency = 13.67 Hz) appears alongside a somewhat different pattern of spiking, relative to the first example. The ISIs are not split across disconnected zones of the histograms (Figure A16.A, *right*), but the presence of bursting episodes is suggested by the high density of very short ISI incidence (28.08% of all ISIs  $\in [2,3]$  ms, compared to a mean of 4.83% for matched, non-bursty synthetic units; see Materials and Methods) and sequential co-occurrence (6.02% of all  $ISI_n, ISI_{n+1}$  pairs  $\in [2,3]$  ms, compared to a mean of 0.25% for the matched synthetic units). Therefore, it is

post-MPTP NHP spike trains: GPi and VLa, thresholded PSDs ( $p_{\text{corr}} < 0.05$ )



**Figure 12.** Points of significant spectral power identified by the shuffling and residuals methods across the population of analyzed spike trains acquired from the parkinsonian non-human primate (NHP). Spike data were sampled and statistically tested according to the same conventions as applied to the two example units in Fig. 5;  $p_{\text{corr}}$  = corrected  $p$  value. (A) Left panel: Points of statistically significant (sig.) power for the full set of GPi (top) and VLa (bottom) uncorrected power spectra. Each heatmap row represents a single power spectrum, bounded within [0, 100] Hz, with each significant point labeled as either a



magenta peak (local maximum within a set of  $\geq 1$  neighboring significant points) or a blue non-peak. Black rectangular outlines mark instances of  $\geq 2$  consecutive significant points falling within the (0,4] Hz delta band (as was used to quantify probable 1/f-like trends). For each region, rows are sorted in order of ascending firing rate (slowest FRs in the highest-positioned rows). Middle panel, right panel: Points of significant power found in the shuffling- and residuals-corrected PSDs. (B) Histograms depicting the count of significant points across frequencies, for the three spectra types and two brain areas.

conceivable that this unit's spiking behavior might also be appropriately described as switching between states of higher FRs (intra-burst) and lower FRs (extra-burst). If so, whether the RP would also be best characterized as varying between the two states is unclear.

These spike pattern observations may again aid interpretation of the PSD results, to some extent. In the uncorrected spectrum, the beta peak sits atop a prominent 1/f-like trend, to which aperiodic burst-associated FR fluctuations could be a contributor [17]. The 1/f-like feature stands out more strongly than any RP distortion that may be present, consistent with the possibility that a very short RP characterizes much of the spike train. The shuffling method does remove the 1/f-like curve, but also removes the beta peak. The residuals method again performs little correction – owing to the 1 ms estimated RP – and therefore retains both the beta peak and 1/f-like trend (with many of the elevated low-frequency points of the PSD also labeled as significant). As alluded to in the Introduction, such 1/f-like features should ideally be modeled and considered an additional component of the baseline against which the statistical significance of candidate oscillatory peaks is determined. Possible strategies for extracting the 1/f-like component – which include both the aforementioned extensions to the residuals method to model bursts, and post-hoc function fits to the corrected PSD – will also be considered further in Chapters 5-7.

#### 4.4.3 GPi and VLa Population Results

Figure 12 illustrates results from the two correction methods for the whole neuronal dataset. The three columns show simplified versions of the spectra (again zoomed in to 0-100 Hz on the  $x$  axis), in which any points that are both significant and labeled as a local spectral peak are labeled in magenta, and all other significant points are marked as blue. Spectra are plotted according to the order of the neuron's observed firing rate, with the lowest-FR units at the top.

The population-level patterns mirror those that we highlight with the individual unit examples. First, we found trends consistent with those that our simulations would predict, with respect to overall incidence of alpha-beta oscillation detections, and the variation of these detections with firing rates. Collapsing across the GPi and VLa units, the residuals method reported that a greater proportion of the total units exhibited significant alpha-beta power, compared to the shuffling method (66/170 versus 37/170 units, McNemar's  $\chi^2 = 19.12, p < 1.3e-5$ ). Moreover, for those units for which alpha-beta band oscillations were reported by at least one method, lower firing rates were predictive of an increased likelihood of residuals-only detection (as compared to shuffling-only or joint method detection;  $\beta_{FR} = -0.089, p < 1.5e-5$ ). Therefore, to the extent that the exclusively residuals-identified power did indeed correspond to true underlying alpha-beta oscillations, the method did fulfill the expectation of enhanced sensitivity, especially in the presence of low firing rates. Note also that the residuals method did identify alpha-beta oscillations over units with a variety of estimated recovery periods (median RP = 3.5 ms, iqr = 6 ms), consistent with the high hit rates seen across the varied RP parameters in the simulations (**Error! Reference source not found.**-Figure A14.).

At the population level, the second pattern – non-oscillatory contributors to spike timing – is most visible as a regular occurrence of 1/f-like effects. In Figure 12, PSDs with strong 1/f-like trends tend to appear as significant points in the lowest frequency bands, starting with the (0,4] Hz delta band, and at times extending well beyond it. Using a simple rule to label likely 1/f-like trends ( $\geq 2$  consecutive significant points within (0,4] Hz; see Chapter 3), we identified a greater incidence of these trends in the residuals- versus shuffling-corrected PSDs (85/170 versus 25/170 units, McNemar's  $\chi^2 = 54.40$ ,  $p < 1.7e-13$ ). This result is consistent with the shuffling method's aggressive removal of the 1/f-like trend in Figure 11B. Also consistent with the Figure 11B unit, we can identify additional cases in Figure 12 for which the residuals method retains a 1/f-like trend that extends into a series of significant alpha-beta points. Therefore, with respect to the above observations of residuals-only alpha-beta detections, an open question does remain regarding the extent to which putative periodic power in these bands would remain significant after removal of the 1/f-like, presumably aperiodic spectral component.

Again consistent with the Figure 11B example, we did observe additional units for which probable bursting may have contributed to 1/f-like trends. For example, for four additional VLa units that met the above delta band rule,  $> 25\%$  of the ISIs fell within 2-3 ms (range: 35.76% - 42.91%), in spite of overall mean FRs of 5.35 Hz - 8.26 Hz.

## 5.0 Two-State Residuals Method: Materials and Methods

### 5.1 Methods Overview

This chapter describes the design and evaluation of the “two-state” residuals method. This exploratory work aimed to address the challenge raised by “bursty” units: To effectively control for these units’ baseline firing properties (i.e., the mean FR and RP), it is necessary to identify (1) the time intervals when the unit resides in the “non-burst” and “burst” states, and (2) the appropriate corrections to apply, to account for state-specific variations in the baseline firing properties. A method that achieves these goals should yield a corrected time series – and ultimately, a corrected power spectrum – that better meets the flat-baseline expectation, compared to the output generated by the one-state residuals method.

The following subsections detail the procedures for (1) simulating bursty units, (2) distinguishing bursty from non-bursty units, and labeling the spike train time bins according to inferred non-burst or burst status, (3) obtaining state-specific RP duration estimates (i.e.,  $\hat{n}_{rNB}$ ,  $\hat{n}_{rB}$ ), and (4) constructing a PPM that incorporates the information obtained through (2) and (3).

Throughout the following, I will use the subscript “ST” to refer to model terms that take on values that depend on the state status. For example, the state-dependent RP duration estimates will be referenced as  $\hat{n}_{rST}$ , where  $ST \in \{ 'NB', 'B' \}$ . Addition of  $(t)$  to this subscript will be used to refer to the value that a parameter is assumed to adopt at a particular time bin, given the inferred burst or nonburst status for that time bin (e.g.,  $\hat{n}_{rST(t)}$  refers to  $\hat{n}_{rB}$  when  $t$  is inferred to fall within a burst state).

## 5.2 Spike Train Simulation: Expansion to Two States

To simulate the states' distinct baseline firing properties, I modified Equation (1) to permit state-specific values for the base spike probability ( $p_{base}$ ), and the duration ( $n_r$ ) and steepness ( $k$ ) of the RP. Therefore, simulation of the spike train followed Equation (5):

$$p_{spk}(t) = \begin{cases} p_{SS}(t) = \mathbf{p}_{baseST(t)} + p_{osc} \cdot \sin(2\pi f_{osc} t \cdot 0.001), & n > n_{rST(t)} \\ p_{RP}(t) = p_{SS}(t) \cdot k_{ST(t)}^{n_{rST(t)+1-n}, & n \leq n_{rST(t)} \end{cases} \quad (5)$$

The simulations assume that the oscillatory drive remains constant throughout the non-burst and burst states (see [69] for evidence consistent with the assumption of independence between bursting and oscillatory phenomena). As in the one-state simulations, I manipulated  $p_{osc}$  indirectly through variation of the modulation index, and specifically through its application to the non-burst base FR:  $m = p_{osc}/p_{baseNB}$ .

The probabilities over the initial spike train states (NB, B) were hard-coded and denoted as  $\pi_{ST}$ . Note that, for parameters that both govern the simulations and are inferred by the Hidden Markov Model (including  $\pi_{ST}$ ; see section 5.4), a hat superscript (e.g.,  $\hat{\pi}_{ST}$ ) will be used to distinguish the inferred from the ground truth values.

To generate a sequence of burst and non-burst dwell times, I departed from the convention from [18] of sampling these inter-transition intervals (ITIs) in units of milliseconds. Instead, I generated ITIs in units of per-visit spike counts (denoted as  $nSpk(j)$ , where  $j$  indexes over all state visitations in a spike train). As such, the millisecond-scale duration of each state visit depended on both the sampled spike count and the ISI durations that arose during the random generation of spikes, according to  $p_{spk}(t)$ . The  $nSpk$ -based approach was motivated by the assumption that the

entries and exits from bursts occur immediately following spikes (as opposed to occurring at some point in the middle of the ISIs). Note that this assumption of spike-triggered transitions is common in the burst detection literature [14]. Moreover, adoption of the time-based transition scheme of [18] could create a biologically questionable scenario in which the RP parameter values change in the middle of a single RP instance (if a spike occurs close to a state boundary).

Each  $nSpk(j)$  was set equal to one plus a random variable drawn from a state-specific geometric distribution:

$$nSpk(j) = 1 + [X \sim G(p_{exitST(t)})] \quad (6)$$

Recall that, given independent trials and a success probability  $p$ , a geometric distribution  $G(p) = p(1-p)^n$  returns the probability that exactly  $n$  failures will be encountered prior to a single success. Therefore, in Eq. (6),  $p_{exitST(t)}$  refers to the probability that the next spike will trigger the transition to the alternative state. To model “doublet” bursts that always comprise a single ISI (such as the Figure 11A example), we can set  $p_{exitB} = 1$ .

As in Chapter 3, spike trains were simulated for  $T$  time bins, where  $T$  was constrained to multiples of 1024 ms. Adherence to this fixed length implied that for the final visited state, fewer than  $nSpk(j)$  spikes might be included in the returned spike train.

To summarize, spike train simulation begins with the sampling of the initial state from  $\pi_{ST}$ , and of a state-specific spike count according to Equation (6). Over a series of subsequent time bins, spikes are generated according to the state-specific parameters for Equation (5), until  $nSpk(1)$  spikes are emitted. On the immediately ensuing time bin, the process switches to the alternate state, and samples a state-appropriate  $nSpk(2)$  value. This process of spike count sampling, spike

generation, and state switching continues through time bin  $T$ . Note that this procedure implies that all state visitations end with a spike (or the end of the simulation).

Note that the spike count-based, geometrically distributed transition scheme does match the assumptions of the Hidden Markov Model that I propose in Section 5.4. In Chapter 7 (Conclusions and Future Directions) I will comment on future priorities for challenging the HMM with datasets that diverge from this model’s core assumptions.

The simulated bursty units (see Sections 5.3 and 5.7.1 for the sets of parameter settings) were matched to “non-bursty” reference units. For a given bursty unit, the reference unit was generated by (1) recording the  $m$  setting and the observed firing rate (FR) for the bursty unit, (2) assuming that  $k = 0$ ,  $n_r = 1$ , and setting  $FR_{ms} = FR/1000$ , setting the non-bursty simulation parameters to  $p_{base} = FR_{ms}/(1-FR_{ms})$ , and  $p_{osc} = p_{base} * m$ . The  $p_{base}$  formula (derived in [53] for  $p_{osc} = 0$  cases, but also reasonably applied in the present  $p_{osc} > 0$  simulations) returns the base firing rate that is expected to produce the target FR, assuming an absolute RP of duration  $n_r = 1$  ms.

### 5.3 Burst detection: Estimation of Initial HMM Settings

To initialize the HMM parameters (prior to their refinement with unsupervised training with the Baum-Welch algorithm; see 5.4), I applied a series of algorithms that estimated the needed distributions from the spike train’s basic ISI statistics.

The following subsections describe the procedures for initializing three sets of HMM priors:

- (1)  $\hat{B}_{ST}$ : The state-conditioned probability distributions over observed ISI durations
- (2)  $\hat{A}$ : The state transition matrix

(3)  $\hat{\pi}$ : The probability distribution over initial burst and non-burst state status

These estimation procedures are illustrated with an example spike train (and its non-bursty reference) that was simulated to approximate the major features of the empirical unit in Figure 11A. For the bursty unit, the state-specific parameters were set as follows (where the non-burst, burst values are presented in the sequence [NB,B]):

(1) base FRs  $p_{baseST} = [19, 550]$  Hz

(2) RP durations  $n_{rST} = [20, 1]$  ms

(3) RP steepness  $k_{ST} = [0.7, 0]$

(4) Initial state probabilities  $\pi = [.98, .02]$

(5) State exit probabilities  $p_{exitST} = [.0435, 1]$

The bursty and reference units shared the following state-general parameters:

(1) Spike train duration ( $T$ ): 120 x 1024 ms

(2) Oscillation frequency ( $f_{osc}$ ): 12 Hz

(3) Modulation index  $m$ : 0.6

### 5.3.1 Preliminary Estimation of State-Specific ISI Distributions

To initialize  $\hat{B}_{ST}$ , I aimed to develop an algorithm that would perform a rough division of the original ISI histogram into component, non-burst and burst distributions. The state-specific ISI distributions could then be directly submitted as the initialized  $\hat{B}_{ST}$  distributions for the Baum-Welch algorithm, which would ideally improve upon the roughness in the initial estimates. An additional design goal was for the output of the histogram-splitting to enable identification of non-bursty units (ideally, by returning an empty object for the burst-specific ISI distribution).



Like the RP estimation algorithm (Chapter 3), the histogram-splitting strategy builds upon concepts introduced by [53]. These authors also sought to identify burst-associated contributions to the spike train, but did not provide an explicit algorithm for achieving this identification (and also assumed uniform RPs for non-bursts and bursts, in contrast to the present work).

The ideas presented by [53] most directly concern the division of the spike hazard function,  $H(ISI)$ , into non-burst and burst components, as opposed to division of the ISI distribution. I adopt these authors' goals of splitting  $H(ISI)$ , and describe an additional step for converting the decomposed hazard function into state-specific ISI distributions.

Recall that, assuming a spike at time 0,  $H(ISI)$  returns the probability that the subsequent spike will occur after  $ISI$  milliseconds, given that this second spike has not occurred yet.  $H(ISI)$  can be computed as  $PDF(ISI)/S(ISI)$ , where  $S(ISI)$  denotes the spike survival function ( $1-CDF(ISI)$ ; i.e., the complementary cumulative density function of the ISIs [6]).

To illustrate how we can detect bursting contributions to  $H(ISI)$ , it is again useful to begin with the simplest hypothetical spike trains, and then expand to more complex cases. To start, consider a non-bursting, non-oscillatory spike train, such as that depicted in Figure A5.. For such a unit, the hazard function would consist of an initial depressed region (due to the RP), followed by an ascent to a roughly flat plateau (indicating the steady state firing probability,  $p_{base}$ ).

Next, consider a non-oscillatory, but bursty spike train, for which the RP remains constant across states, but  $p_{baseB}$  greatly exceeds  $p_{baseNB}$ . In this case, the initial dip is first followed by a peak, and then resolution to the flat region. The peak represents the short ISIs that occur more frequently than expected, given  $p_{baseNB}$ . [53] proposed that one could identify burst-specific contributions to  $H(ISI)$  by isolating an early post-RP peak that rises above a neighboring flat (albeit likely noisy) line.

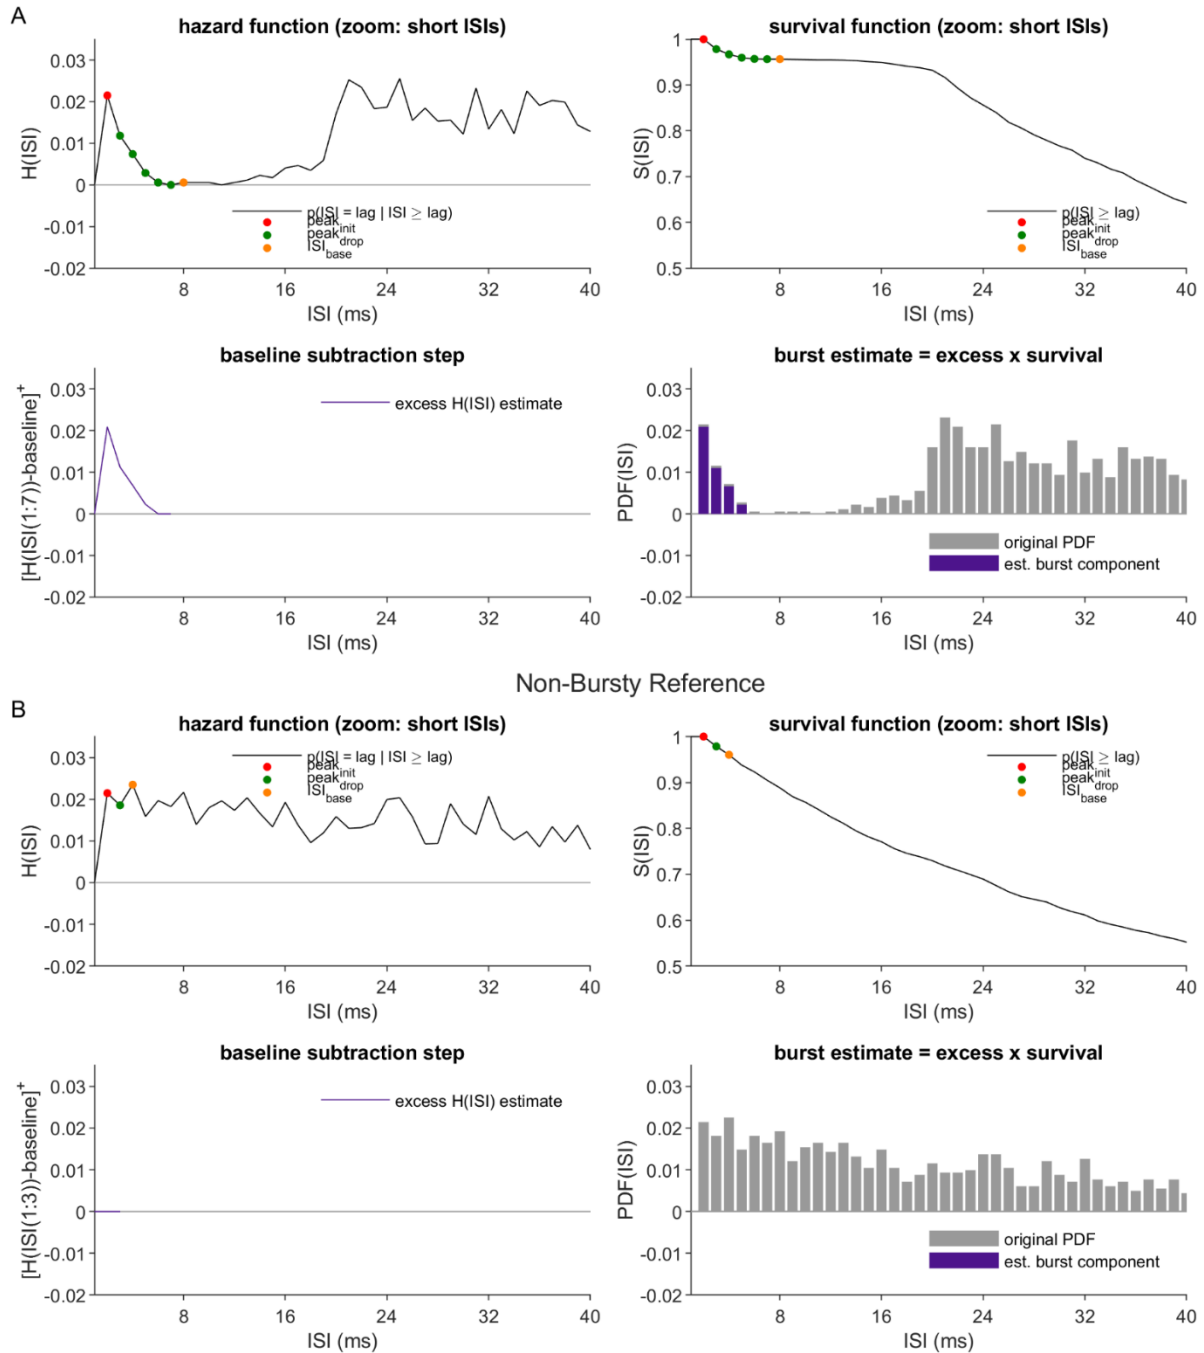
Unfortunately, this simple peak-isolation strategy is disrupted by the introduction of state-dependent RPs. In the present synthetic example (Figure 13A), the very long non-burst RP creates a second, post-peak depression in  $H(ISI)$ , followed by a rise to the non-burst steady state region. The algorithm that I propose is sensitive to this possibility, and seeks a peak that rises above a baseline defined over only a neighboring point in the hazard function (instead of a sequence of several neighboring points).

I do assume that other sources of spike train complexity (e.g., more modest differences in RP duration, and oscillatory trends) exert sufficiently little impact on  $H(ISI)$  to significantly disrupt the peak and baseline identification schemes I describe below. Modest inaccuracies in the divided hazard function (and the resulting state-specific ISI distribution estimates) are tolerable, since the end objective of this procedure is to produce rough initial estimates only.

Given the above, the algorithm for splitting  $H(ISI)$  is illustrated in Figure 13A, and proceeded as follows:

- (1) Given a delta vector (that is, the 1 kHz, binary time series where “1” marks a spike-containing time bin), compute both the ISI distribution and  $H(ISI)$ , with time bins spanning the ms lags in  $[1, \max(ISI)]$ .
- (2) Using the original RP estimation algorithm (Section 3.4.3.1), identify the ISI corresponding to the start of the post-RP region (i.e.,  $\hat{n}_r + 1$ ). In Figure 13A, this point on the  $H(ISI)$  plot appears as a red dot (labeled as “ $peak_{ini}$ ”, since this point marks the start of a candidate burst-related peak).
- (3) Identify the remainder of the candidate peak (“ $peak_{drop}$ ”) as a consecutive run of ISIs for which  $H(ISI)$  is decreasing; i.e.,  $dH(ISI)/dISI < 0$ . In Figure 13A, these points on the  $H(ISI)$  plot appear as green dots.

Synthetic Approximation to Fig. 11A Unit: Decomposition into Two ISI Distributions  
Bursty Spike Train



**Figure 13. Demonstration of the procedure for separation of an ISI distribution into burst- and non-burst components. (A)**

Application to an example spike train simulated to approximate the empirical unit from Figure 11A. This synthetic unit was generated to express two features: a non-burst RP duration that greatly exceeds that of the burst RP, and bursts characterized by a single, and typically very short ISI. The separation procedure identifies burst-related activity through a peak in the hazard function (following from [53]), and multiplies the isolated peak by the survival function, to obtain the bursting contribution to the original

ISI PDF. See Section 5.3 for the full set of simulation parameters and abbreviation definitions. (B) Application of the separation procedure to a non-bursty reference unit, simulated to match the overall spike rate and oscillation properties of the bursty unit (see 5.3 for details). Consistent with the non-bursty status of this unit, no short-ISI peak in the hazard function is identified.

(4) Let “peak<sub>end</sub>” refer to the final point in the peak<sub>drop</sub> region. Set the ISI at peak<sub>end</sub> + 1 to “ $ISI_{base}$ ” (orange in Figure 13).  $H(ISI_{base})$  will determine the reference value against which the candidate peak region is compared.

(5) Let  $H(ISI_{RP+peak})$  refer to the vector of values returned by evaluating  $H(ISI)$  over the range  $[1, peak_{end}]$ . Compute the “excess” area under the hazard function as  $[H(ISI_{RP+peak}) - H(ISI_{base})]^+$ , where  $[ ]^+$  denotes zero-rectification.

The “excess” area (see Figure 13A, 2<sup>nd</sup> row, left panel), and the area that remains after its extraction, are assumed to reflect the burst and non-burst components of the hazard function ( $subH(ISI)_B, subH(ISI)_{NB}$ ), respectively. Multiplication by  $S(ISI)$  converts these  $subH(ISI)_{ST}$  to the corresponding decomposition of the original ISI PDF, which I will denote as  $subPDF_{ST}$  (see Figure 13A, 2<sup>nd</sup> row, right panel). To generate the initial estimates for  $\hat{B}_{ST}$ , we can normalize these two ISI components to state-specific PDFs (“ $PDF_{ST}$ ”), which can directly serve as the initialized state-specific observation matrices for the HMM (that is,  $\hat{B}_{ST} = PDF_{ST}$ ). Note that the HMM described below (Sections 5.4-5.5) assumes that the ISI duration is a categorical, rather than continuous variable, thereby enabling this direct use of  $PDF_{ST}$  for the initialization (rather than a function that has been fit to it, such as a sum of Gaussians [50]). Section 5.4 will provide further information on the rationale behind this choice of categorical representation.

Figure 13B depicts the output of the  $\hat{B}_{ST}$ -initialization procedure for the non-bursty reference unit. Since the spike train arose from a single set of  $p_{base}$ ,  $n_r$ , and  $k$  settings,  $H(ISI)$  does not exhibit a clear peak region. Moreover, because  $H(ISI_{base})$  exceeds all points in  $H(ISI_{RP+peak})$ , the

decomposition procedure returns an empty vector for  $subH(ISI)_B$ . Given such an outcome, the unit is labeled as “non-bursty”, and dropped from the downstream steps described below. One may note that, for some non-bursty units,  $\geq 1$  points in  $H(ISI_{RP+peak})$  will exceed  $H(ISI_{base})$  by chance, resulting in inaccurate output of a nonzero  $subH(ISI)_B$  estimate. Although the proposed procedure would process such a unit with the full burst detection pipeline, the very modest mass assigned to  $subPDF_B$  in these cases will influence the subsequent pipeline steps (see below) in a manner that will help to counteract the risk of excessive false-positive burst labeling.

### 5.3.2 Preliminary Estimation of Initial State Distribution

The estimated prior probabilities that an ISI would arise from a burst or non-burst state ( $p_B$ ,  $p_{NB}$ ) were set to  $\text{sum}(subPDF_B)$ , and  $1-p_B$ , respectively. The distribution over start states ( $\hat{\pi}$ ) was populated with these priors.

### 5.3.3 Preliminary Estimation of State Transition Matrix

I detail below the procedure for initializing  $\hat{A}$  (and see Appendix B.1, Figure B1. for a schematic depiction). To preview the approach at a high level: As stated in section 5.2, the HMM I propose here shares the simulations’ assumption that the number of spikes that “fail” to trigger a state exit (equivalent to the count of ISIs minus one) follows a geometric distribution, defined by the single parameter  $p_{exitST}$  (see Eq. 6). Therefore, this preliminary estimation step attempts to (1) obtain rough initial estimates (as compared to that returned by an HMM) of the location of the burst and non-burst segments in the ISI time series, (2) tabulate the counts of exit “failures” for each segment

assigned to the assumed burst and non-burst states, respectively, and (3) estimate the  $p_{exitST}$  values by fitting geometric distributions to this failure count data. The resulting  $\hat{p}_{exitST}$  estimates provide sufficient information to populate the off-diagonal elements of  $\hat{A}$ , as these entries indicate the likelihood of entering the alternate state on the next timestep, given the state label for the present timestep. The main diagonal entries are constrained to be the complements of the  $\hat{p}_{exitST}$  entries.

The estimation algorithm proceeded as follows:

- (1) Given a delta vector, compute the corresponding sequence of ISIs (“ $ISI$ ”).
- (2) Let  $P(ISI / ST) = PDF_{ST}$ . Combining these state-conditioned observation probabilities with the appropriate state priors ( $p_{ST}$ ), use Bayes’ rule to compute the posterior likelihoods of the states, over the set of possible ISIs:  $P(ST | ISI) = P(ISI | ST) * p_{ST} / P(ISI)$ .
- (3) For each state, evaluate  $P(ST / ISI)$  at each of the ISIs in  $ISI$ . Let  $w_{ISI_{ST}}$  ( $w$  = “weighted”) refer to the vector outputs of this evaluation step. Therefore, the  $i$ th entry of  $w_{ISI_{ST}}$  represents an estimated posterior probability that  $ISI_i$  was sampled from state  $ST$ , given solely the information determined through step (2).
- (4) Steps (4)-(8) segment  $ISI$  into candidate burst and non-burst intervals. First, identify all contiguous runs of nonzero entries in  $w_{ISI_B}$  (which will typically alternate between runs of nonzero and zero entries, since  $P(B / ISI)$  will inherit the many zero entries from  $P(ISI / B)$ ).
- (5) The candidate burst segments identified in (4) will vary with respect to the  $P(B / ISI)$  weights that they include. Moreover, the total number of ISIs that contribute to these segments may exceed the count expected on the basis of the estimated  $p_B$  value. To prepare to narrow down the list of candidate segments to a high-weight subset, first compute the mean value of the  $w_{ISI_B}$  entries in each segment (“ $w_{ISI_{BSeg}}$ ”).

- (6) Sort  $wISI_{BSeg}$  in order of descending weight (“ $swISI_{BSeg}$ ”)
- (7) Retain only the candidate burst segments associated with the first  $n$  elements in  $swISI_{BSeg}$ , such that  $n$  maximizes the number of segments that may be included, without exceeding the total count of ISIs anticipated by  $\hat{p}_B$ . I will refer to this list of retained burst segments as  $Seg_B$ .
- (8) Label all ISIs not in  $Seg_B$  (that is, not within the retained  $n$  segments of top weight from  $swISI_{BSeg}$ ) as non-burst segments ( $Seg_{NB}$ ). In other words, for all ISIs that do not belong to a segment that is a member of  $Seg_B$ , new segments are formed for each of the runs of  $\geq 1$  ISI that are not interrupted by a  $Seg_B$  segment, and this collection of new segments is assigned to the  $Seg_{NB}$  set.
- (9) For each state, compute the list of ISI counts pertaining to each segment ( $CT_{ST}$ ). Subtracting 1 from each count in this list ( $CT_{ST} - \vec{1}$ ) returns a list of counts of “failures” encountered before a spike triggers a transition to the alternate state (e.g., for a state visit consisting of one ISI, the spikes defining the ISI mark the entry to and exits from the state, implying 0 failures).
- (10) For each state, find the maximum likelihood fit (here, with Matlab *mle*) of the geometric distribution ( $p(1-p)^n$ ) that models the data present in  $CT_{ST} - \vec{1}$ . The  $\hat{p}$  values for the distribution fits can serve as the estimates for  $\hat{p}_{exitST}$ : the likelihood that the next spike encountered in a state will trigger the exit from that state.
- (11) Construct  $\hat{A}$  as follows:

$$\begin{bmatrix} 1 - \hat{p}_{exitNB} & \hat{p}_{exitNB} \\ \hat{p}_{exitB} & 1 - \hat{p}_{exitB} \end{bmatrix}$$

## 5.4 Burst detection: Refinement of HMM Parameters

For the next step of the burst detection process, the ISI vector, and the initial  $\hat{B}_{ST}$ ,  $\hat{\pi}$ , and  $\hat{A}$  estimates, are submitted to the Baum-Welch algorithm (also known as the “forward-backward” algorithm). Several references provide helpful reviews of this common procedure [48, 49]. To briefly review, Baum-Welch is an Expectation Maximization (EM)-based algorithm. Therefore, like more broadly familiar EM algorithms (e.g., K Means), Baum-Welch iteratively cycles through Expectation (E) and Maximization (M) Steps until reaching convergence (potentially on a local maximum, as mentioned in Chapter 2). In the E Step, the current working version of the model’s parameters, in combination with the HMM assumptions listed in 2.9.1, and the input data, are used to update quantities related to the inferred probabilities for the latent state assignments at each timestep  $t$ . In the M Step, a set of rules are applied to the quantities determined during the E Step, to return updated  $\hat{B}_{ST}$ ,  $\hat{\pi}$ , and  $\hat{A}$  parameters that improve the model’s likelihood of having generated the observed data.

With reference to [70], I implemented the Baum-Welch algorithm in Matlab. Computations were performed over log-converted variables to address the challenges raised by very small probabilities [71]. The algorithm terminated according to a log likelihood (LL)-based criterion, such that the working version of the HMM parameters was returned once  $LL(n) - LL(n-1) < 10^{-4}$ . Although this convergence criterion did rely on a hard-coded free parameter (which this pipeline has been designed to largely avoid), note that adjustment away from the  $10^{-4}$  tolerance setting does not meaningfully impact the final results.

Finally, as noted in 5.3.1, the observation probabilities  $\hat{B}_{ST}$  were assumed to reflect a categorical, rather than continuous variable, such that the Baum-Welch update rules were applied



to each of the ms-width ISI bins independently. This option was chosen after pilot testing of an alternative procedure that treated the ISI durations as a continuous variable. Specifically, the  $\text{PDF}_{\text{ST}}$  were fit with Gaussian Mixture Models, and the GMM parameters served as the inputs to and refined outputs of the Baum-Welch step (following the update rules from [70]). This GMM-HMM ultimately produced poorer burst detection accuracy, relative to the categorical approach.

### 5.5 Burst detection: HMM-Based Decoding of ISI States

To apply the Baum Welch output to infer a sequence of latent states, I implemented the Viterbi algorithm in Matlab (again adopting the logarithm approach from [71]). Similar to Baum-Welch, the Viterbi algorithm combines the input data, the model parameters, and the HMM assumptions to draw inferences regarding the likelihoods of the possible latent states (and here, the possible nonburst versus burst states underlying each observed ISI). Specifically, Viterbi returns the sequence of “hard” (i.e., not probabilistic) state assignments that maximize the likelihood of the input data. Therefore, in the context of the current problem, Viterbi returns a sequence of hard ‘NB’ or ‘B’ classifications for each in the series of input ISIs.

As a brief recap (and see also [72]), we can summarize the Viterbi algorithm as a recursive procedure coupled with memoization (i.e., maintenance of computed values in a table). In the context of the present problem (and adopting the  $k$  index from Chapter 2, to refer to samples instead of timepoints), Viterbi constructs a  $2$  (# of states)  $\times$   $K$  (# of ISIs) matrix  $V$ , where each element  $V_{ST,k}$  stores the maximum possible likelihood of having encountered ISIs  $1:k$ , under the assumption that the underlying process resides in state  $ST$  at sample  $k$ . For  $k = 1$ , the elements in  $V_{ST,1}$  are

initialized according to the corresponding entries in  $\hat{\pi}$  (the estimated probabilities of each start state). At  $k > 1$  (and with  $ST'$  referring to the assumed state at the  $k-1$  sample),  $V_{ST,k}$  is updated as:

$$V_{ST,k} = \operatorname{argmax}_{ST'} [V_{ST',k-1} \cdot \hat{A}_{ST',ST} \cdot \hat{B}_{ST,ISI_k}]$$

In other words, for each possible prior state  $ST'$  occurring at  $k-1$ , the Viterbi algorithm computes the product of three probabilities: the maximum observation likelihood estimate that was stored for the prior state ( $V_{ST',k-1}$ ), the probability of entering state  $ST$  on  $k$ , assuming that state  $ST'$  occurred at  $k-1$ , and the probability of observing  $ISI_k$ , assuming that state  $ST$  occurs at  $k$ . The maximum over these products is stored as  $V_{ST,k}$ . The algorithm also stores, for each state and  $k > 1$ , a “backpointer”  $bp_{ST,k}$ , which indicates the identity of the prior state  $ST'$  that satisfied the  $\operatorname{argmax}$  term. At the final sample ( $k = K$ ), a likelihood-maximizing path through the V matrix is identified by (1) selecting the row  $ST$  in column  $K$  with the maximal  $V$  value, and (2) reconstructing the path by tracing the backpointers from path end to path start, beginning with  $bp_{ST,K}$  for the  $ST$  value selected in step (1). The resulting sequence provides the state labels that Viterbi returns.

The proposed “two-state” PPM (see Section 5.6) requires a sequence  $z$  of state labels that are aligned with the original millisecond-width time bins, such that  $z(t) \in \{ 'NB', 'B' \}$ . In translating the ISI-aligned state labels to the time bin-aligned case, I adopted the convention of filling in the time bins that preceded and included the first spike, and that followed the last spike, with ‘NB’ labels.

## 5.6 Two-State Residuals-corrected PSD

Once the burst labels are available, the two-state residuals method entails three steps: (1) estimation of state-specific recovery period durations ( $[\hat{n}_{rNB}, \hat{n}_{rB}]$ ), (2) fitting a PPM to the spike train, which estimates state-specific base FR and  $t - s_*(t)$  effects (up to a bound set by  $\hat{n}_{rST}$ ), and (3) applying spectral analysis to the residuals of the PPM fit.

### 5.6.1 Estimation of Two-State Recovery Period Durations

Relative to the one-state RP estimation algorithm (Section 3.4.3.1), the two-state procedure requires three modifications. First, the two-state algorithm operates on state-specific ISI distributions, here corresponding to  $PDF_{ST}$ . The second and third modifications address challenges created by the burst ISIs. Specifically, as became apparent from the synthetic  $PDF_B$  generated for this work, the ISI that should be identified as  $\hat{n}_{rB} + 1$  may commonly be the maximum lag with a nonzero ISI count (which, in the original algorithm, marked the end of the ISI histogram). It is problematic for the ISI histogram to end at this point, since the original algorithm identifies  $\hat{n}_r + 1$  by finding the point of the first local  $\Delta D$  maximum. Therefore, to ensure that the algorithm would attempt a curve fit that is left-anchored to the correct  $\hat{n}_{rB} + 1$  point, it was necessary to zero-pad the ISI histogram to the right by one additional ISI. In addition, to allow the  $\Delta D$  for such a curve fit to stand out as a local maximum, it was necessary to add one bin of zero-padding (again, to the right) to the vector of  $\Delta D$  estimates output by the curve-fitting process.

### 5.6.2 Two-State Bounded Last-Spike Point Process Model

In contrast to the one-state PPM, the two-state PPM learns coefficients for state-specific intercept and RP-related terms. Let  $\alpha_0$  and  $\alpha_1$  refer to binary representations of the non-burst and burst state labels, such that  $\alpha_0(t) = 1$  for all  $t$  where  $z(t) = \text{'NB'}$ , and  $\alpha_1 = \neg \alpha_0$ <sup>8</sup>. The two-state PPM takes on the form

$$\log \lambda(t | H_t, \beta) = \sum_{k=0}^1 \alpha_k(t) \left[ \beta_{0k} + \sum_{j=1}^{\hat{n}_r(k)} \beta_{jk} \delta_{t-j, s^*(t)} \right] \quad (7)$$

### 5.6.3 Spectral Analysis of the Two-State Model Residuals

Once a solution is obtained for (7), spectral analysis proceeds in the same manner as described for the one-state method. The raw residuals are calculated as:

$$r(t) = I(t) - e^{\left[ \sum_{k=0}^1 \alpha_k(t) \left[ \hat{\beta}_{0k} + \sum_{j=1}^{\hat{n}_r(k)} \hat{\beta}_{jk} \delta_{t-j, s^*(t)} \right] \right]} \quad (8)$$

The  $r(t)$  output was prepended with sufficient zero-padding to ensure a final vector of length  $T$ .

---

<sup>8</sup> Here,  $\neg$  represents the bitwise NOT operator (that is,  $\alpha_1(t) = 1 - \alpha_0(t)$  for all elements of the time series).

## **5.7 Preliminary Evaluation of the Two-State Residuals Method**

The shuffling, and the one- and two-state residuals methods (referenced as 1S- and 2S-residuals, respectively) were compared with respect to their ability to support accurate oscillation detection. As in the original residuals versus shuffling comparison, this three-method comparison entailed (1) simulation of synthetic spike train datasets, (2) generation of the methods' corrected PSDs, with tabulation of oscillation hit and false alarm rates for these PSDs, and (3) statistical comparison of the methods' classification of true and false oscillations. Because this thesis reports only a preliminary evaluation of the 2S-residuals method, the synthetic datasets, and the assessments performed upon the methods' output for these datasets, is much more limited than that reported in Chapters 3-4.

### **5.7.1 Two-State Synthetic Spike Train Test Sets**

Bursty spike trains, and their matched non-bursty units, were generated using three sets of parameters. The first set, listed in section 5.3, aimed to approximate the qualitative features seen in the empirical unit from Figure 11A. The second and third parameter sets (Table 1) aimed to either approximate the Figure 11B unit (left column) or challenge the ISI histogram-splitting procedure with a moderate-length nonburst RP (which contributes a complicating positive trend to the burst-related region of the histogram).

**Table 1. Parameters for Simulation of Two Additional Sets of Bursty Units.** Bracket entries indicate [nonburst, burst] settings. “ST” = state (referring to state-dependent settings). NB = Non-burst. All remaining abbreviations follow from those used in Figures 2-3.

Table 1. Parameters for Simulation of Two Additional Sets of Bursty Reference Units		
	Fig. A11B Match	Moderate NB RP
base FRs $p_{\text{baseST}}$	[18, 500] Hz	[18, 500] Hz
RP durations $n_{\text{rST}}$	[1, 1] ms	[9, 1] ms
RP steepness $k_{\text{ST}}$	[0, 0]	[0.7, 0]
Initial state probabilities $\pi$	[0.95, 05]	[0.95, 05]
State exit probabilities $p_{\text{exitST}}$	[0833, 0.3333]	[0833, 0.3333]
Spike train duration T:	120 x 1024 ms	120 x 1024 ms
Oscillation frequency $f_{\text{osc}}$ :	12 Hz	12 Hz
Modulation index m:	0.6	0.6

In both of the above cases, the non-bursty reference units shared the  $T$ ,  $f_{\text{osc}}$ , and  $m$  settings with their associated bursty units, with the remaining parameters set according to the specifications described in Section 5.2. For each of the three sets of parameter settings, 20 bursty and 20 non-bursty reference units were generated.

### 5.7.2 Evaluation of Oscillation Detection Outcomes

The methods for identifying statistically significant points in the PSDs, and for classifying PSDs as hit- and FA-containing, were the same as those reported in Chapter 3. The tests for significant PSD points were limited to a single  $\alpha_c$  value of .05, and no subsampling was performed.

### 5.7.3 Evaluation of Burst Detection Outcomes

Given the importance of burst detection to the overall 2S-residuals method, a limited set of tests assessed the performance of this step of the pipeline. Towards this end, I calculated two sets of measures. The first and second sets focused on the classification of bursty versus non-bursty units, and the classification of burst- and non-burst ISIs (for bursty units), respectively. For the first set, I summarized the overall accuracy, and the hit and false alarm rates for the discrimination between bursty and non-bursty spike trains. Here, hit rate = # of correct bursty unit guesses / # of bursty units, and false alarm rate = # of incorrect bursty unit guesses / # of non-bursty units. For the second set (applied to the bursty units specifically), I referred to the ground truth, ISI-aligned state labels to compute the overall accuracy, and the hit and FA rates for the detection of burst-associated ISIs (with similar hit/FA definitions applying at the level of the bursty units' ISIs, as were applied to bursty versus non-bursty unit discrimination). For the overall accuracy measure, I compared the ISI-aligned output of the Viterbi algorithm with the output of a dominant class-guessing strategy (i.e., entire vectors of 'NB' labels).

## **6.0 Two-State Residuals Method: Results**

The following sections summarize the performance of the two-state residuals pipeline on synthetic bursty and non-bursty units generated with the three parameter sets of interest. For each parameter set, I begin by illustrating the operation of the method on a representative example generated with each of the parameter sets.

Subsequently, I report analyses over the full datasets of simulated units, in which I compared the two-state (2-S), one-state (1-S), and shuffling methods' corrected PSDs. These analyses confirm that expansion from one to two states improves the residuals method's ability to support accurate oscillation detection in bursty spike trains, without a noticeable detriment to oscillation detection in non-bursty spike trains.

### **6.1 Example #1: Long Non-Burst Recovery Period and Two-Spike Bursts**

The initial parameter set was used to generate a synthetic bursty unit that approximated the empirical unit from Figure 11A. A key recreated feature was the clear separation of the presumed non-burst and burst zones in the ISI histogram. As referenced in Chapter 5, Figure 13 illustrates the operation of the histogram-splitting algorithm on this histogram, and also on the ISIs for a non-bursty control unit (matched to the bursty unit's mean FR, and oscillatory frequency and modulation strength; see 5.2). Since the algorithm did not assign any non-zero mass to the non-



bursty unit's "burst" PDF, this unit was correctly identified as non-bursty, and not subjected to additional processing.

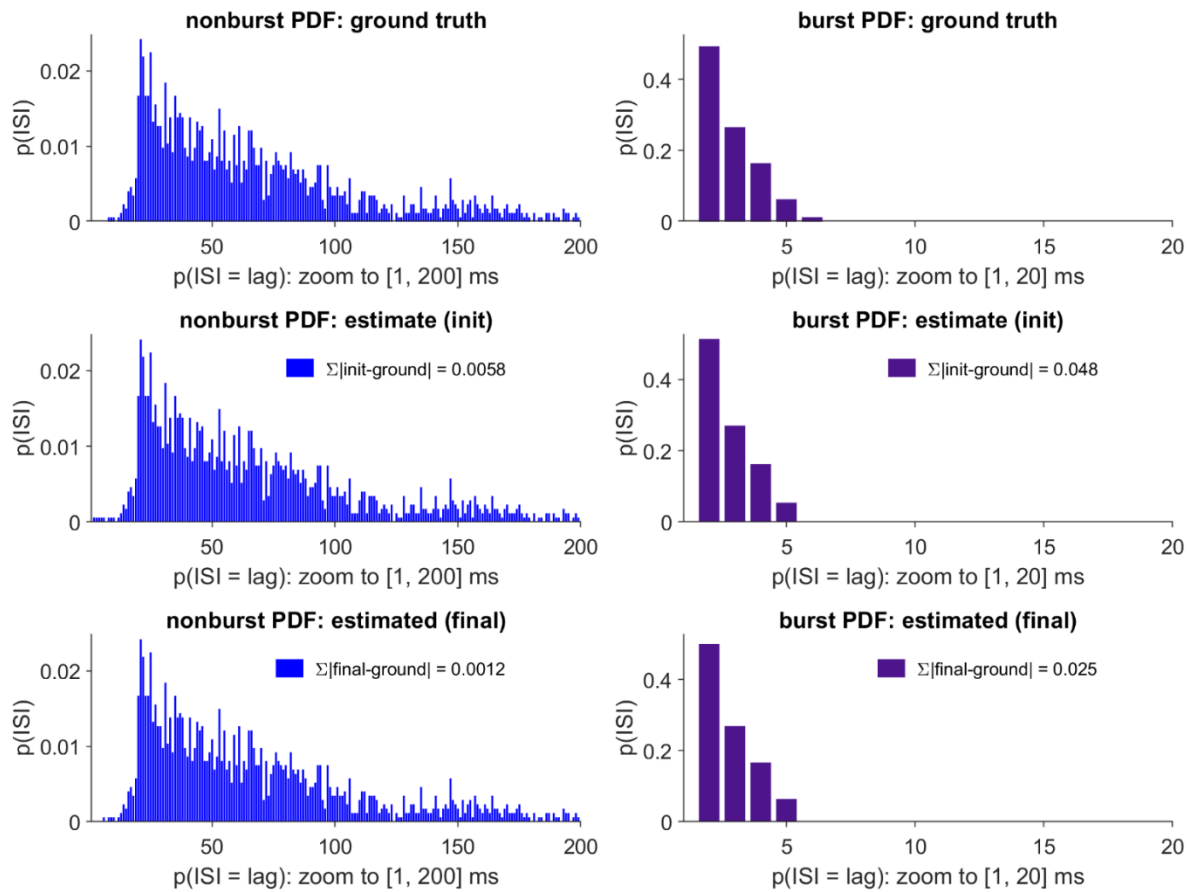
For the bursty unit, Figure 14 shows the state-specific PDF estimates generated by the histogram-splitting procedure (2<sup>nd</sup> row), the refinement of these PDF estimates by the Baum-Welch algorithm (3<sup>rd</sup> row), and the target, ground truth PDFs (1<sup>st</sup> row).

In this case, the Baum Welch step resulted in only very small changes in the final ISI probabilities (maximum single bin decrease = 2.5%). The final estimates appear generally similar to the ground truth estimates, with only a few small perceptible differences from the ground truth PDF (mostly involving the addition or omission of very low-height probability values).

The HMM-based burst detection step made use of this estimated PDF, in addition to estimates of the start state and state transition probabilities, to infer the non-burst versus burst status of the ISIs in the spike train. The only labeling errors consisted of a small count of missed burst ISIs (accuracy, and burst hit and FA rates = 99.94%, 98.73%, 0%, compared to 95.65% accuracy for a fixed 'NB' guess). In generation of the two-state PPM, these labels were combined with the estimated RP durations for the two states ( $\hat{n}_{rNB}, \hat{n}_{rB} = [19, 1]$  ms, as compared to the ground truth  $n_{rNB}, n_{rB} = [20, 1]$  ms).

The power spectra for the two-state and original one-state PPMs' residuals, and the shuffling method are shown in Figure 15. Similar to the approximated empirical unit (Figure 11A), this simulated unit gave rise to a beta-frequency oscillation that was detected by the shuffling method, but missed by the 1S-residuals method (which under-corrected for the distortion generated by the long-duration non-burst RPs). The 2S-residuals method did effectively flatten out the distortion, resulting in successful oscillation detection (and also omission of the false alarm points generated by the 1S method).

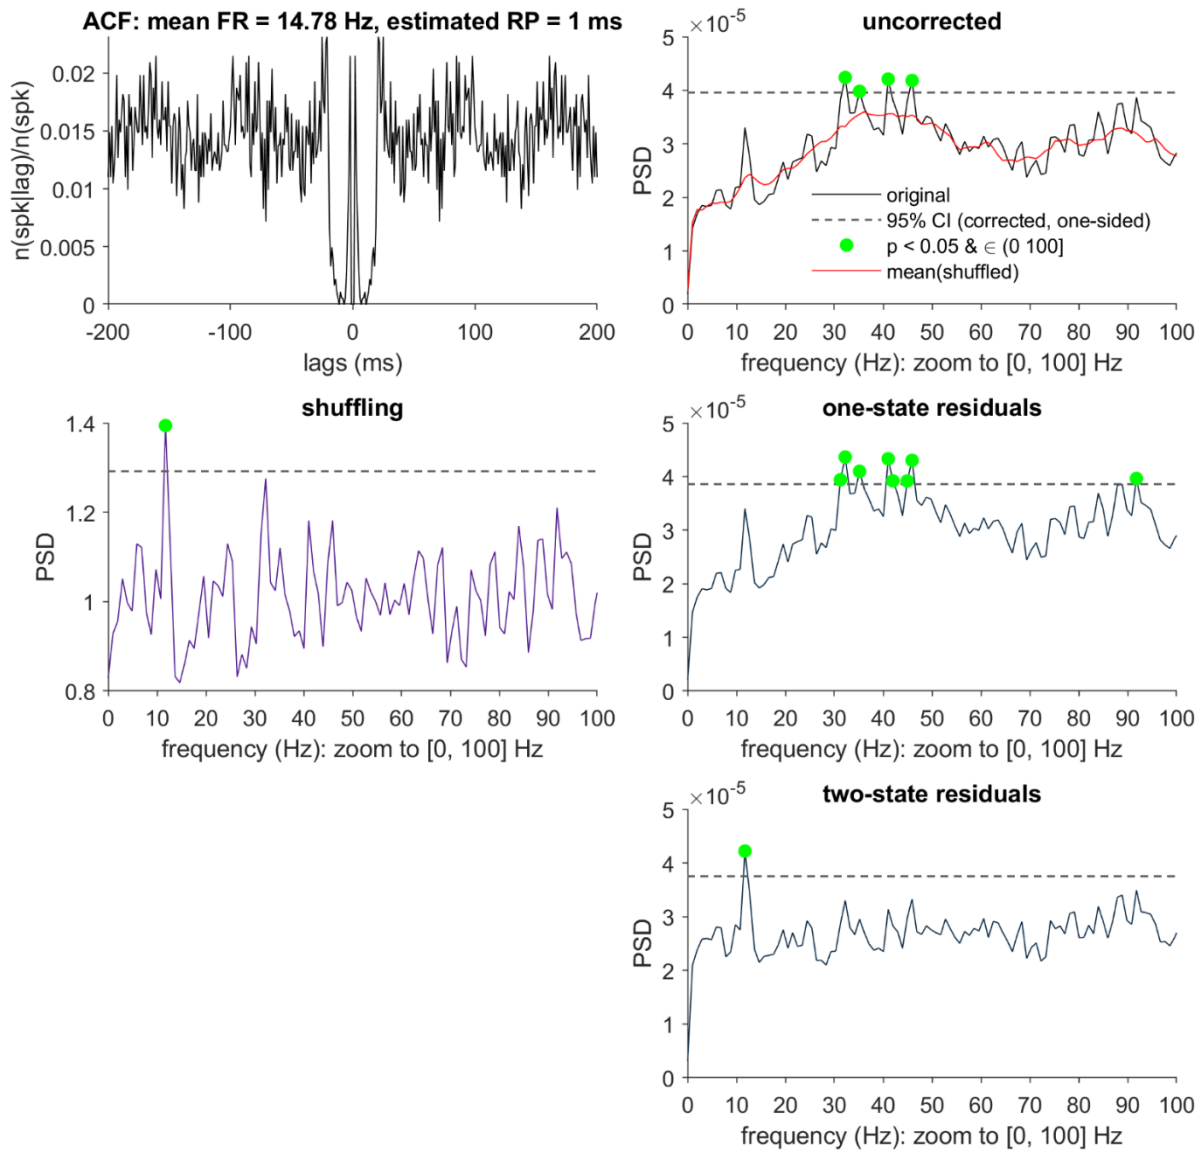
Synthetic Approximation to Fig. A11A Unit (Bursty): Comparison of Ground Truth and Estimated State PDFs



**Figure 14.** True and estimated ISI probabilities for the synthetic approximation of the empirical unit illustrated in Figure 11A. Ground truth (top), initial, pre-Baum Welch (init) and post-Baum Welch (final) estimates for these distributions.  $\Sigma|init-ground|$  and  $\Sigma|final-ground|$  refer to the sum of the absolute bin-wise differences between the initial estimated and ground truth distributions, and final estimated and ground truth distributions, respectively.

Synthetic Approximation to Fig. A11A Unit (Bursty): Comparison of Corrected PSDs

100 Hz scale



**Figure 15. Comparison of the shuffling output, and the one-state (1S), and two-state (2S) residuals output for the unit simulated to match the empirical unit from Figure 11A. Formatting conventions and abbreviations follow from those used in Figure 10.**

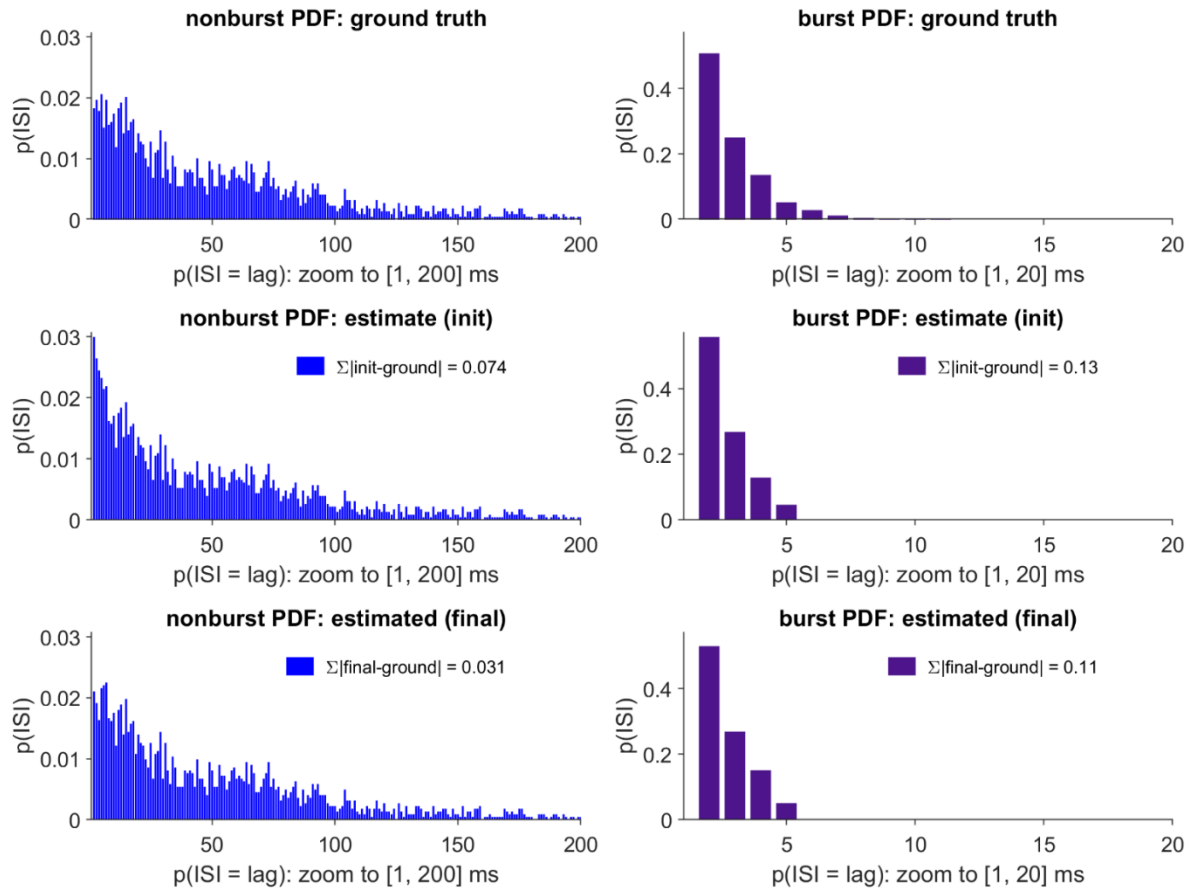
## 6.2 Example #2: Matching Two-State RPs and a Prominent 1/f-like Trend

The second parameter set was used to generate a synthetic bursty unit that approximated the empirical unit from Figure 11B. As it was ambiguous whether this unit was appropriately modeled as arising from two distinct RPs (see Figure A16.B), the simulation was conducted under the conservative assumption that only the base FR differed between non-burst and burst states. Therefore, the synthetic bursty unit presented a situation in which considerable overlap existed between the two states' ISI distributions, for the shortest ISI durations.

The stages of ISI-splitting process, for both the simulated bursty and non-bursty referenced units, are shown in Figure B2.. Although the initial split did appear to allocate excessive mass to the shortest ISIs in  $\text{PDF}_{\text{NB}}$  (i.e., the high overlap zone), application of Baum Welch aided in removing much of this excess (Figure 16). The final allocation generally resembles the ground truth PDFs. Similar to the previous example, the remaining visible discrepancies principally concern very short histogram bins (e.g., the omission of the longest-duration ISI mass from the  $\text{PDF}_{\text{B}}$ ).

For the inferred ISI state labels, the accuracy, hit, and false alarm rates were 95.21%, 83.27%, and 1.78%, respectively (with 79.90% accuracy produced by an all-NB guess). Therefore, to the extent that the HMM did commit errors, these predominantly consisted of non-burst labeling of true “burst” ISIs. The  $\hat{n}_{r_{\text{NB}}}$ ,  $\hat{n}_{r_{\text{B}}}$  estimates matched the ground truth values of [1,1] ms.

Synthetic Approximation to Fig. A11B Unit (Bursty): Comparison of Ground Truth and Estimated State PDFs



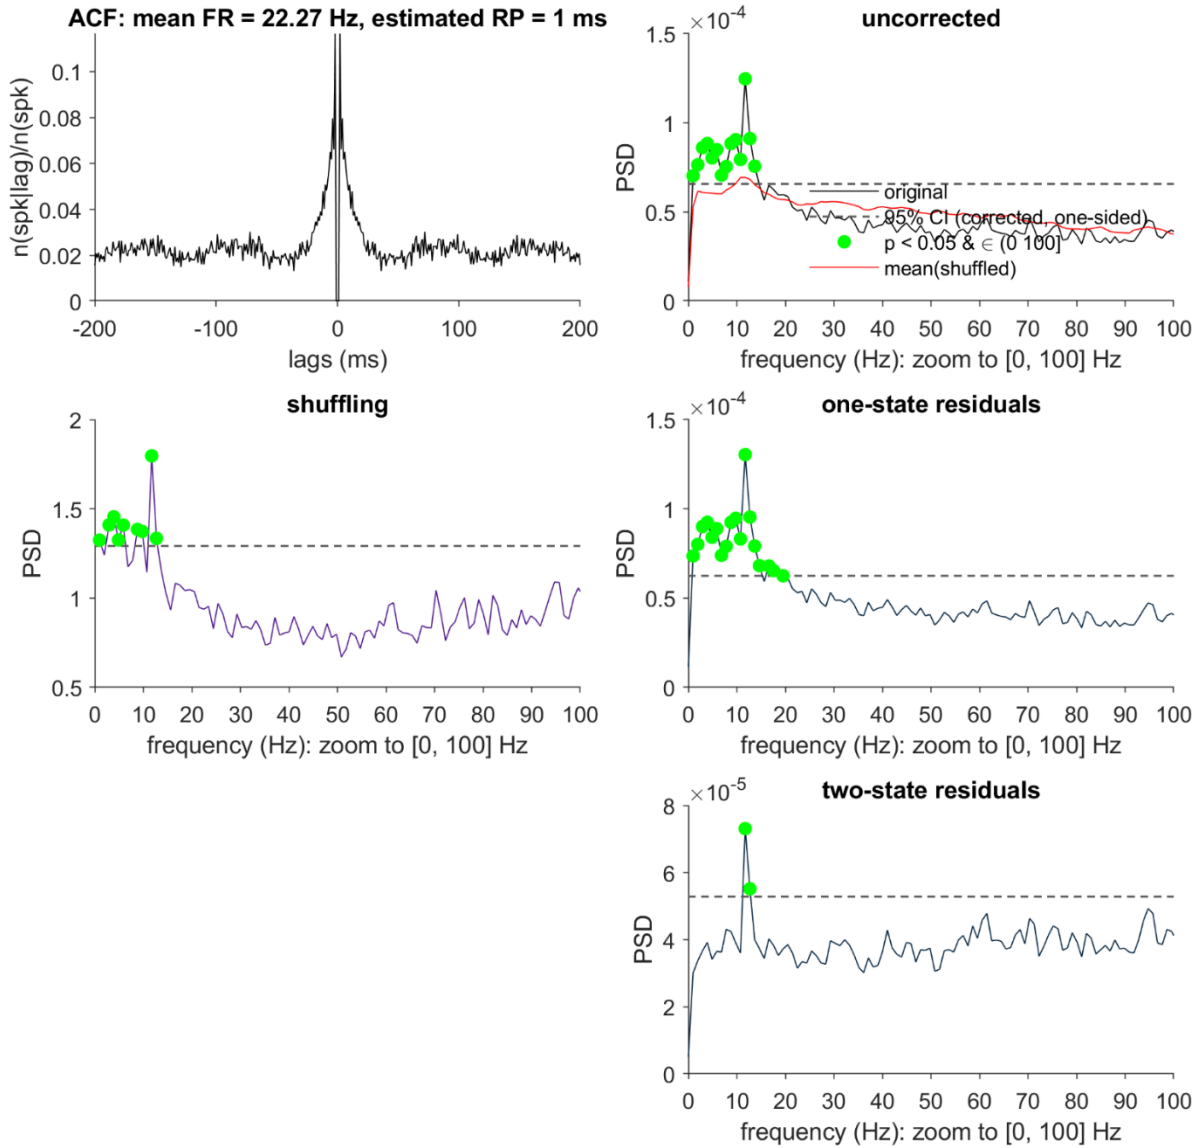
**Figure 16.** True and estimated ISI probabilities for the synthetic approximation of the empirical unit illustrated in Figure 11B. Ground truth (top), initial, pre-Baum Welch (init) and post-Baum Welch estimates for these distributions. See Figure 14 for legend descriptions.

The corrected PSD produced by the 2S-method (Figure 17), indicates that the benefits conferred by this approach remained robust, in spite of the errors in ISI labeling. Unlike the 1S-residuals output (and the shuffling output, in this case), the 2S output exhibited a successful flattening of the prominent  $1/f$ -like trend that contributed to the uncorrected PSD. Consequently, the 2S-method selectively flagged beta points as significant, while omitting the low-frequency false positive points present in the 1S case.

For the non-bursty reference unit, a very limited non-zero allocation to  $subPDF_B$  appears at  $\text{ISI} = 2$  ms (amounting to .18% of the height of the original ISI bar; see Figure B2.). This result

Synthetic Approximation to Fig. A11B Unit (Bursty): Comparison of Corrected PSDs

100 Hz scale



**Figure 17.** Comparison of the shuffling, one-state (1S) residuals output and two state (2S) residuals output for the unit simulated to match the empirical unit from Figure 11B. Formatting conventions and abbreviations follow from those used in Figure 10.

implies that the unit is incorrectly labeled as “bursty”, but the risk of false-burst labels is mitigated by the downstream steps’ consideration of the trivial  $subPDF_B$  contribution. In the ISI sequence, the HMM labeled only labeled the single, initial ISI as burst-generated. Accordingly, the differences between the PSDs output by the 2S- and 1S- methods were negligible (Figure B3.).

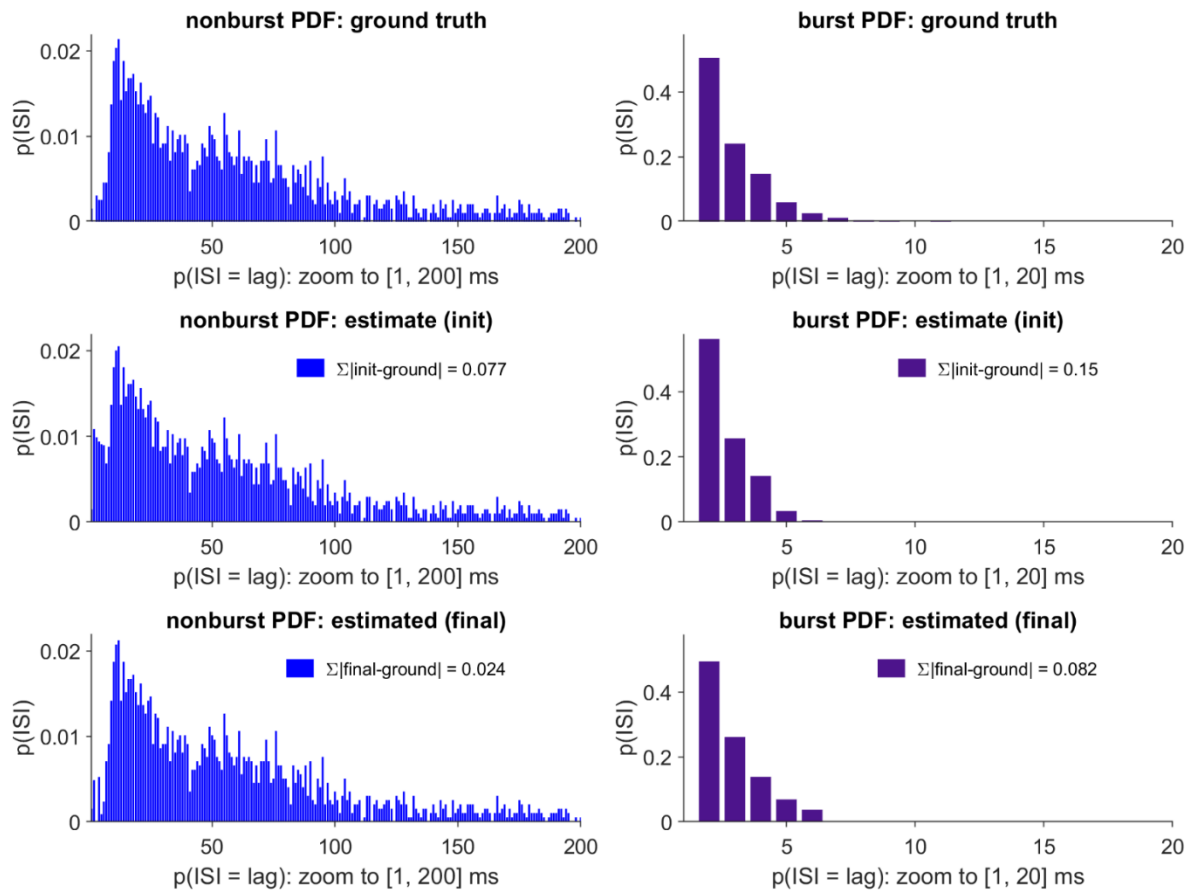
### 6.3 Example #3: Moderate Non-burst RP and Prominent 1/f-like Trend

Implicit in the current histogram-splitting algorithm is an assumption that the burst-related peak rises above a region of the hazard function that would be flat in its absence (following the initial RP dip). For the previous two parameter sets, any RP-generated departures from this assumption were expected to be minimal (due to the imposition of either uniform RPs, or a non-burst RP that remains flat over the ISIs associated with the burst peak). For the third parameter set, the non-burst states of the bursty units were generated with a moderate-duration RP that was expected to contribute to a rising trend to the same region of the hazard function in which the burst-related peak was expected to occur (while the RP remained short for the burst states).

Figure B4 shows the results of the histogram-splitting process for the simulated bursty and reference units. For the non-bursty unit, an empty object is returned for  $subPDF_B$ , and so this unit is rejected for further processing. For the bursty unit, Figure 18 shows the ground, initial, and final state-specific PDF estimates. As in the previous, uniform-RP example, the initial  $PDF_{NB}$  estimate reflects an overestimation of the probability of very short ISIs in this state, which is later corrected, to a great extent, by the Baum Welch procedure. This step also makes small but noticeable adjustments in the shape of  $PDF_B$ .

For the ISI-level labels, the burst detection accuracy, hit, and false alarm rates were 98.49%, 95%, and .76%, respectively (with accuracy = 82.39% for the all-NB guess). The RP estimates  $\hat{n}_{rNB}$ ,  $\hat{n}_{rB}$  matched the ground truth values of [9,1] ms. Consistent with the generally reasonable

Synthetic Unit, Moderate-Length NB RP (Bursty): Comparison of Ground Truth and Estimated State PDFs



**Figure 18. True and estimated ISI probabilities for a synthetic approximation for a synthetic unit with opposite-direction RP and burst effects on the shortest lags in the hazard function.** Ground truth (top), initial, pre-Baum Welch (init) and post-Baum Welch estimates for these distributions. See Figure 14 for legend descriptions.

accuracy of the above estimates, the corrected PSD produced by the 2S-residuals method (Figure 19) achieved the expected flattening of the spectrum, with removal of the low frequency false alarm points that were present in the uncorrected, shuffling, and 1S-residuals PSDs.



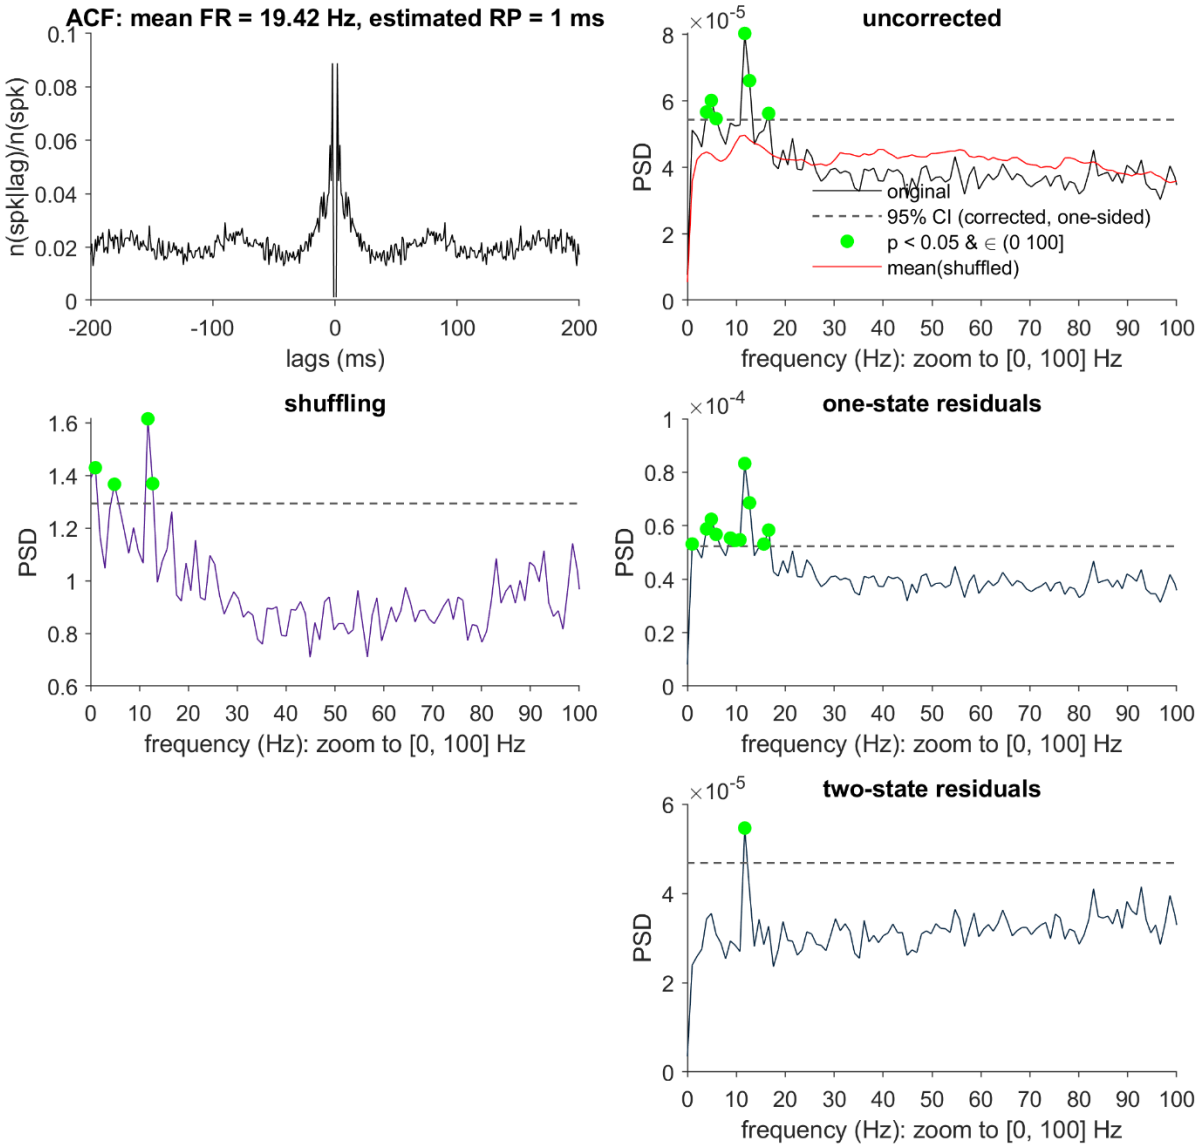
## 6.4 Summaries over Multiple Simulations

As an initial assessment of the generality of the above results, I performed a sequence of analyses on small datasets of multiple spike trains ( $n = 20$  per combination of 3 parameter categories  $\times$  2 “burstiness” status levels). The initial analyses (restricted to the true “bursty” units, see Table 2) compared the rates of hit- and FA-containing PSDs generated by the shuffling and the one- and two-state shuffling methods, at the single  $\alpha = 0.05$  level.

These summary statistics confirm that the observations from the single simulation run examples remain largely consistent across runs. In the case of two highly divergent RPs (Set 1), the 2S-residuals method maintains perfect detection of the oscillation and avoidance of false alarms, in contrast to the occasional missed detection and persistent FA commission for the 1S method. For the two parameter sets with strong  $1/f$ -like trends (Sets 2-3), the 2S-residuals method does, for one instance, miss an oscillation that the shuffling and 1S-residuals methods report, but this small decrease in hit rate is accompanied by a complete avoidance of false alarms (which appear in 100% of the shuffling and 1S-residuals PSDs). Therefore, at least for these parameter settings, the 2S-residuals method appears to improve upon the 1S method’s performance, in such a way as to further improve upon the residuals versus shuffling advantage.

Synthetic Unit with a Moderate-Length Non-burst RP: Comparison of Corrected PSDs

100 Hz scale



**Figure 19.** Comparison of the shuffling, one-state (1S) residuals output and two state (2S) residuals output for the unit simulated with a moderate-length non-burst RP duration. Formatting conventions and abbreviations follow from those used in Figure 10.

**Table 2. Rates of hit-containing and false alarm (FA) containing PSDs for the “bursty” units.** Results are shown for the shuffling, one-state (1S) residuals, and two-state (2S) residuals methods, for the “bursty” units generated with each of three parameter settings (n = 20 iterations per setting).

Table 2. Rates of Hit and False Alarm-Containing PSDs for 3 Correction Methods and 3 Parameter Settings			
Row	Set 1	Set 2	Set 3
shuffling: hit rate	1	0.9	1
shuffling: FA rate	0.15	0.9	1
1S-residuals: hit rate	0.7	0.9	1
1S-residuals: FA rate	1	0.9	1
2S-residuals: hit rate	1	0.9	0.95
2S-residuals: FA rate	0	0	0

Subsequently, we can ask whether the burst detection stage of the two-state pipeline successfully distinguished between the bursty and non-bursty control units (Table 3):

**Table 3. Performance of “bursty” versus “nonbursty” classification for the two-state residuals method.**

Columns report, for the three parameter sets, the likelihood that the burst detection stage of the two state residuals method would label a bursty unit as “bursty” (top row), or a non-bursty unit as “bursty” (second row), and the maximum percentage of the original ISI distribution’s probability mass that was incorrectly assigned to a “burst” distribution, across the 20 non-bursty units tested within each parameter set.

Table 3. Rates of correct and false positive detection of bursty units, and misassignment of non-burst probability mass as part of a burst			
	Set 1	Set 2	Set 3
p(bursty bursty)	100%	100%	100%
p(bursty not-bursty)	55%	60%	55%
non-bursty units: max % of PDF to PDF <sub>B</sub>	1%	2%	0.86%

This table shows that “bursty” units were always identified as such. The false alarm rates for misidentifying non-bursty units as “bursty” were high, but it is useful to remember that the criterion is very lax for such a judgment, and that the actual fraction of the original distribution that is allocated to an inferred “bursty” PDF plays the most important role in the subsequent burst labeling

step. Over all the non-bursty units simulated across the three parameter sets, at most 2% of the entire original ISI distribution was attributed to the estimated “burst” distribution.

Finally, we can summarize the burst labeling performance at the level of the individual ISIs within the bursty units (Table 4):

**Table 4. For the three parameter sets, the accuracy, hit (second row), and false alarm (third row) rates for the labeling of ISIs within individual spike trains.** Values are reported as means +/- SD. The final row reports the accuracy of a uniform guess of “non-burst” for all ISIs.

	Set 1	Set 2	Set 3
mean accuracy	99.67% +/- 0.23	94.68% +/- 0.85	98.04% +/- 0.49
p("burst"   burst ISI)	99.58% +/- 0.79	88.66% +/- 3.95	97.60% +/- 0.61
p("burst"   nonburst ISI)	0.33% +/- 0.24	3.80% +/- 1.36	1.85% +/- 0.64
accuracy, all non-burst guess	95.67% +/- 0.38	79.97% +/- 1.61	80.36% +/- 1.53

Again, these dataset-level trends mirror those seen in the individual examples. Performance is generally reasonable. The most noticeable area for improvement concerns a tendency for occasionally missed “burst” ISI labels in the Set 2 cases (characterized by the high ISI distribution overlap).

## **7.0 Conclusions and Future Directions**

### **7.1 Publishing Note**

Note that a subset of the text in this Chapter has been adapted from a manuscript that has been posted to a preprint server [19] and submitted to a peer-reviewed journal.

### **7.2 Chapter Overview**

To recap, my principal objective was to develop accurate computational methods for identifying oscillations in spike trains – and, by extension, point process time series with similar properties. A key obstacle to this goal was the concurrent shaping of spike trains by aperiodic phenomena, which can introduce trends that can be difficult to disentangle from those introduced by oscillations. In particular, the aperiodic features can generate elevated regions in the power spectral density (PSD) of the spike train, which a naïve oscillation-detection algorithm might mistakenly attribute to a truly periodic origin. Therefore, I focused on developing methods for removing the distortion that specific non-oscillatory phenomena can contribute to the PSD. Successful removal would allow for oscillations to be identified through the use of statistical tests that assume that significant oscillations will appear as sufficiently large deviations from an otherwise flat PSD.

The principal motivator for this work was the problem of correcting distortion attributable to the neuronal recovery period (and related phenomena in other domains), and the bulk of this thesis

focuses on a novel correction method that addresses limitations in earlier approaches to solving this specific problem. A follow-up, exploratory effort concerned the expansion of this method to simultaneously correct for the effects of “bursting” on the PSD.

In the following sections, I will first summarize the major predicted findings of the initial, RP correction effort (i.e., the “one-state” residuals method), and a less anticipated finding, concerning the behavior of the method in the presence of high firing rates. Subsequently, I will recap the contributions of the expanded “two-state” residuals method towards addressing the bursting-related challenges that the one-state method left open. Finally, I will touch upon priorities for future research, including further improvements of the residuals method, and additional evaluations that would benefit development of the two-state method (which was subjected to only preliminary evaluation here).

### **7.3 One-State Residuals Method**

In Chapter 2, I reviewed an established ISI shuffling method [7] that approaches the recovery period problem by extracting the spectral component explained by the spike train’s ISI distribution alone. Because the ISI distribution contains information about not only the recovery period, but also oscillations, the shuffling method may hinder detection of relatively subtle oscillatory spiking, especially in slowly spiking units. As an alternative, I developed a regression-based “residuals” method that also models and removes temporal structure related to the ISIs, but only those ISIs that fall within a historical bound determined by the estimated duration of the recovery period (RP) function. I compared residuals-corrected versus shuffling-corrected PSDs to determine whether

the former enhanced sensitivity to oscillations, with this effect predicted to be greatest for low FR spike trains.

### 7.3.1 Predicted Findings

The two methods were initially evaluated on synthetic spike trains, generated with varying settings for firing rate, oscillatory frequency and modulation strength, and recording duration, and with the RP represented as an exponentially-recovering modulation of spike probability. Collapsing over these variable settings, I found that residuals correction enabled more accurate classification of true versus false positive points in the power spectra. More detailed examination of the varied parameters, and of their impact on the methods' relative hit and false alarm rates, revealed three general trends. First, and as anticipated, residuals correction primarily benefited from enhanced sensitivity to the ground truth oscillations, especially when firing rates were low, relative to the oscillation frequency. Second, two other predictors of strong relative residuals performance (moderate oscillation modulation ( $m$ ); short spike train duration ( $T$ )) suggest that this method may be most useful for identifying subtle oscillatory effects that reside close to a threshold level of detectability. Third, the largely sensitivity-driven benefit was gained, for most parameter settings, without exaggeration of false positive oscillation detections. On a limited subset of high FR, strongly oscillating spike trains, both the residuals and shuffling methods did tend to produce PSDs with very low-power FA peaks; I propose strategies for addressing these FAs farther below.

The largely favorable results for the residuals method generalized across variable recovery period durations (3, 4, 9, and 18 ms), and shapes (i.e., absolute or exponentially-rising, with variable steepness). This persistence of successful detection across RP characteristics was

accompanied by relatively accurate estimation of the ground truth RP durations by the ISI distribution-fitting algorithm, with estimation errors rarely exceeding 1-2 ms.

In the real GPi and VLa data – keeping in mind the caveat concerning the unknowability of the ground truth – I did observe additional evidence consistent with our expectations. Focusing on the alpha-beta frequency range that was anticipated in the parkinsonian NHP under consideration, I found that the residuals method reported significant oscillations for a greater fraction of the units than the shuffling method did. This result was anticipated on the basis of the prevalence of low-to-moderate FRs in this dataset (especially amongst the VLa units), and follow-up tests confirmed the relevance of low spike rates for predicting residuals-only oscillation detection. Although we cannot presently determine whether the residuals-flagged oscillations are “real”, it is conceivable that future empirical work could aid in providing corroborating evidence. For example, confidence that an oscillatory drive governs a unit’s spiking could be bolstered by the discovery of matching oscillatory trends in the unit’s synaptic input dynamics (e.g., as tracked by high resolution biosensors of the relevant neurotransmitters [45]). One might also anticipate that any significant spike oscillations may correspond to similar oscillations in a unit’s membrane voltage. Consistent with this expectation, a recent study in mice [46] observed that striatal cholinergic interneurons that exhibited robust delta rhythms in their spiking (as estimated in the time domain with an ISI-based heuristic) also showed strong delta rhythms in subthreshold membrane voltage (as estimated with wavelets).

A subset of the GPi and VLa spike trains exhibited bursting and 1/f-like trends that the one-state residuals method was not designed to address. Consequently, the corresponding residuals PSDs should be interpreted with care. Section 7.4 discusses the potential for the two-state residuals method to address these patterns.



### 7.3.2 Effect of false alarms with fast spiking and strong modulation

When operating on synthetic inputs generated with high base FRs ( $p_{base}$ ) and strong oscillatory modulation ( $m$ ), both shuffling and residuals correction produced spectra with elevated false alarm rates (Figure 9). This trend worsened with increasing spike train duration ( $T$ ), and with variation in the “true” oscillation frequency ( $f_{osc}$ ). The nature of the  $f_{osc}$  effect depended on the method in question. For the lowest frequencies (e.g., 7 Hz; Figure A8.), residuals correction commonly yielded false positives, but the shuffling FA rate remained mostly low. As  $f_{osc}$  approached the highest tested settings (e.g., 32 Hz; Figure A9.), shuffling FAs become extremely frequent, while the residuals FA rate moderated somewhat, such that the relative specificity of the two methods flipped. Under the FA-promoting conditions, both methods reported false peaks that tended to fall within the gamma band, and were characterized by very low power that barely exceeded the statistical significance threshold (**Error! Reference source not found.A**). The distinctive features of these false alarms suggest that a simple PSD postprocessing step (e.g., imposition of a more stringent significance and/or effect size threshold) may help minimize their appearance in future analyses.

Future development of the residuals method should seek to prevent these elevated false alarms. In the fast spiking, strong modulation cases, a key challenge concerns the heightened negative correlation between the effect that we are attempting to model (the recovery period) and the unmodeled oscillation effect. This correlation engenders an “omitted variable bias” [73] in RP function estimation. Visualization of representative spike trains (**Error! Reference source not found.B**) illustrates the core difficulty: For a strongly oscillating spike train, the spikes and the recovery period regressors will principally cluster under the positive phases of the oscillation

cycle. Consequently, the RP regressors are associated not only with the recent occurrence of single spikes, but also with an elevation of the background firing rate above its overall mean. The result is an underestimation of the RP's suppressive effect, which in turn leads to incomplete removal of the RP distortion effect from the corrected PSD.

Note that aggravated correlation between the RP and oscillation effects also clarifies the tendency for the false alarms to be most prominent for the lowest oscillation frequency cases. As illustrated in

Figure A11.B, the correlation between the recovery period and oscillatory components of the spike-generating rate function is more negative for the 7 Hz as opposed to the 32 Hz example. In the more slowly oscillating case, the broadened positive phases of the oscillation often span the full duration of the recovery periods (as opposed to just the first few milliseconds, as in the faster oscillation case), thereby further complicating the ability of the regression model to estimate a suppressive effect for the RP.

Stevenson [73] noted that omitted variable biases might be reduced through introduction of variables that may capture suspected unmodeled effects (here, oscillations of unknown frequency), but also acknowledged that such strategies can be assumption-intensive and nontrivial. Such variables are present in those PPMs, discussed in Chapter 2, that estimate spike rhythms entirely in the time domain (e.g., [11-13]). Therefore, when a researcher is considering the tradeoffs between adoption of either one of the time-domain methods, or one of the alternatives (i.e., shuffling or the proposed residuals method), the extent to which omitted variable bias is a concern should be kept in mind.

## 7.4 Two-State Residuals Method

The expansion from the one- to the two-state residuals method was intended to address two hypothesized contributors to the apparently suboptimal performance of the one-state method on select empirical units. The first hypothesis, concerning the VLa unit presented in Figure 11A, was that this unit alternated between bursting and non-bursting states, for which the RP duration differed substantially. The one-state residuals model cannot accommodate distinct burst and non-burst recovery periods, thereby implying that the RP-related variance would be incompletely addressed in such a two-RP case. In the VLa example, the model estimated and removed very short RP effects, which appeared to appropriately describe bursting states in the spike train, but was highly mismatched to the apparently long RPs that occurred throughout the remainder of the spike train. Consequently, this residuals-processed PSD showed trends consistent with distortion under-correction.

The second hypothesis concerned the VLa unit in Figure 11B. For this unit, I observed a prominent,  $1/f$ -like trend in the original, uncorrected PSD (and also in the one-state residuals output), and also patterns in the ISIs that were suggestive of bursting. Therefore, it seemed reasonable to speculate that correction of the highly divergent firing rate states that bursting creates (an established contributor to  $1/f$ -like trends [17]) might aid in reducing the  $1/f$ -like trend in the power spectra (which, like the RP, creates a problematic broadband distortion that disrupts the flat-baseline assumptions).

Based on preliminary evaluations, the two-state method appears to demonstrate some potential to help address concerns such as those described above. The complete method comprises a multi-stage pipeline, starting with a burst detection step that both labels burst-related intervals in the

spike train, and returns non-burst and burst-associated ISI distributions. The ISI distributions may be used to estimate state-specific RP durations, and this information, in combination with the spike train-aligned state labels, may be used to construct a suitable two-state PPM.

When tested on synthetic units designed to mimic the Figure 11A example, the various stages of the pipeline appeared to operate effectively: The burst detection step showed some success in identifying the non-burst and burst states, and the RP durations that were characteristic of those states. Consequently, the two-state method returned PSDs for which the under-correction problem was successfully addressed.

Similarly, the pipeline appeared to work effectively on two synthetic units that exhibited visible  $1/f$ -like trends (including one designed to approximate Figure 11B, and a slightly modified version of this unit). The burst detection method achieved reasonable identification of the highly overlapping non-burst and burst contributions to the ISI histogram, in addition to sufficiently accurate burst-labeling and RP identification to support an effective two-state PPM. The residuals of these regression models demonstrated flattening of the  $1/f$ -like trend, and removal of the accompanying false alarm points.

The above said, it is important to recognize two limitations of the current method. First, note that bursts represent just one of the many diverse underlying dynamics that may give rise to  $1/f$ -like trends [74]. Therefore, in real, empirically acquired data, adoption of the 2S-residuals method will be insufficient, on its own, to account for all such effects. A complementary correction procedure – which is already commonplace in research using continuous signals – involves the standard postprocessing step of removing a fitted  $1/f^\chi$  function from the PSD (where  $\chi$  is a standard exponent parameter that is added to accommodate variability in curve steepness; a “knee” component may also be added, to account for an extra bend). The routine introduced by [63] offers

one option for achieving a robust curve fit to the PSD. This method additionally fits a series of Gaussians to any narrowband spectral peaks, and those Gaussians that sufficiently exceed the aperiodic baseline are inferred to represent likely oscillations.

A second limitation concerns the current status of the burst detection performance. Although the histogram-splitting and burst-labeling steps showed some signs of success, performance is certainly not yet at ceiling. Therefore, efforts to further improve this stage of the pipeline (especially for units with relatively high degrees of nonburst and burst ISI overlap, as seemed to present the greatest challenge to the current approach) would be helpful.

## **7.5 Areas for Future Work**

Beyond the above, two additional areas stand out as priorities for future work.

### **7.5.1 Further Evaluation of the Burst Detection Step**

First, more extensive evaluation and visualization of the burst detection stage of the 2S-residuals method would be helpful. As discussed in Chapter 2, burst detection is a challenge that has attracted considerable interest in its own right, thereby adding to the motivation for assessing the currently proposed method's performance. Expanded evaluation could take a few different forms. To start, it will be important to test the detection algorithm on a wide variety of synthetic and empirical datasets. The synthetic datasets should include those that are simulated through frameworks that were developed by other authors, for other purposes, and with blindness to both

the approach proposed here, and the previous work that informed it [53]. Evaluation on such a broad test set would help clarify the extent to which this new detector is useful for capturing the full range of bursting patterns that are of interest in neuroscience, in contrast to solely those that evolve in a manner that weds very closely to the proposed detector's current working assumptions.

Future evaluation should also expand upon the set of metrics used to judge the detector's performance. For example, the evaluation here only reported summaries of ISI-by-ISI comparisons of the estimated versus ground truth labels. It may be useful to complement these measures with more detailed information about how the errors were distributed in time (e.g., scattered as isolated errors across the time series, clustered within long runs or around the time of state transitions, and so on).

Finally, the proposed detector should be compared to other, established burst detection methods. The newly proposed method does offer a means of burst detection that requires no user-input parameters (apart from the tolerance threshold for Baum Welch convergence). Therefore, it would be useful to learn if its performance is sufficiently strong to justify its use as a potential replacement to those algorithms for which the configuration demand is high.

### **7.5.2 Adaptation to Time-Frequency Analysis**

In their present forms, both the one- and two-state residuals methods assume the presence of stable oscillatory components. Adaptation of the residuals framework to the analysis of time-varying oscillations may be feasible, depending on two factors: the unit's spike rate, and the stability of its RP function. Vigorous spiking and a stable RP present the optimal scenario. In this case, one could reasonably apply the one-state PPM strategy, and submit the residuals time series

to the preferred time-frequency analysis method (e.g., short-time Fourier Transform (STFT), Wavelet Transform, or bandpass filtering followed by Hilbert Transform).

If spiking is slow, then insufficient data may be available to detect transient oscillations, even with adequate RP correction. Simulations may aid in estimating the minimal spike rate required for sensitivity to any hypothesized effects.

If the RP function is likely to vary over time, the pattern of the suspected variation should inform the adaptation approach. If one assumes that the RP varies across transitions between two latent, discrete states, but remains stable within states, then it may be reasonable to first apply the two-state PPM (as we have described for bursting) and then submit the residuals to time-frequency analysis. Alternatively, if one assumes both fast spiking and RP stability within each segment of an STFT analysis, one might conceivably execute the full one-state residuals routine (RP estimation/PPM/PSD) on each segment independently. More challenging data features (e.g., complicated RP variation patterns) will require the development of methods that go beyond the basic approaches that we have laid out here.

If time-frequency analyses are attempted, then appropriate postprocessing may need to be applied to additionally extract any remaining  $1/f$ -like trends. For example, Wilson et al. [75] recently proposed a time-resolved method for aperiodic and periodic trend decomposition, which extends upon the stationarity-assuming method of Donoghue et al. [63].

## 7.6 Concluding Remarks

Throughout this thesis, I have described a number of existing methods for addressing the non-oscillatory trends that might disrupt the process of seeking oscillations in spike trains. A single “best” method is unlikely to exist. Each method offers a distinct set of pros and cons, and researchers must identify the option that best matches their analysis goals. The residuals method is an especially useful option when a frequency domain representation is required, spike rates are low, and when other factors, such as modest modulation strength and short recording duration, challenge oscillation detection. In the implementation I present here, the method draws upon an accessible and flexible GLM framework. Moreover, the generation of a corrected time series output – and not only corrected oscillation estimates – expands the set of analyses that the underlying framework might ultimately support, including common time-frequency analyses.



## Appendix A Supporting Information, Figures and Tables for the One-State Residuals

### Method

#### Appendix A.1 Shuffling Correction by Division: Rationale and Assumption Evaluation

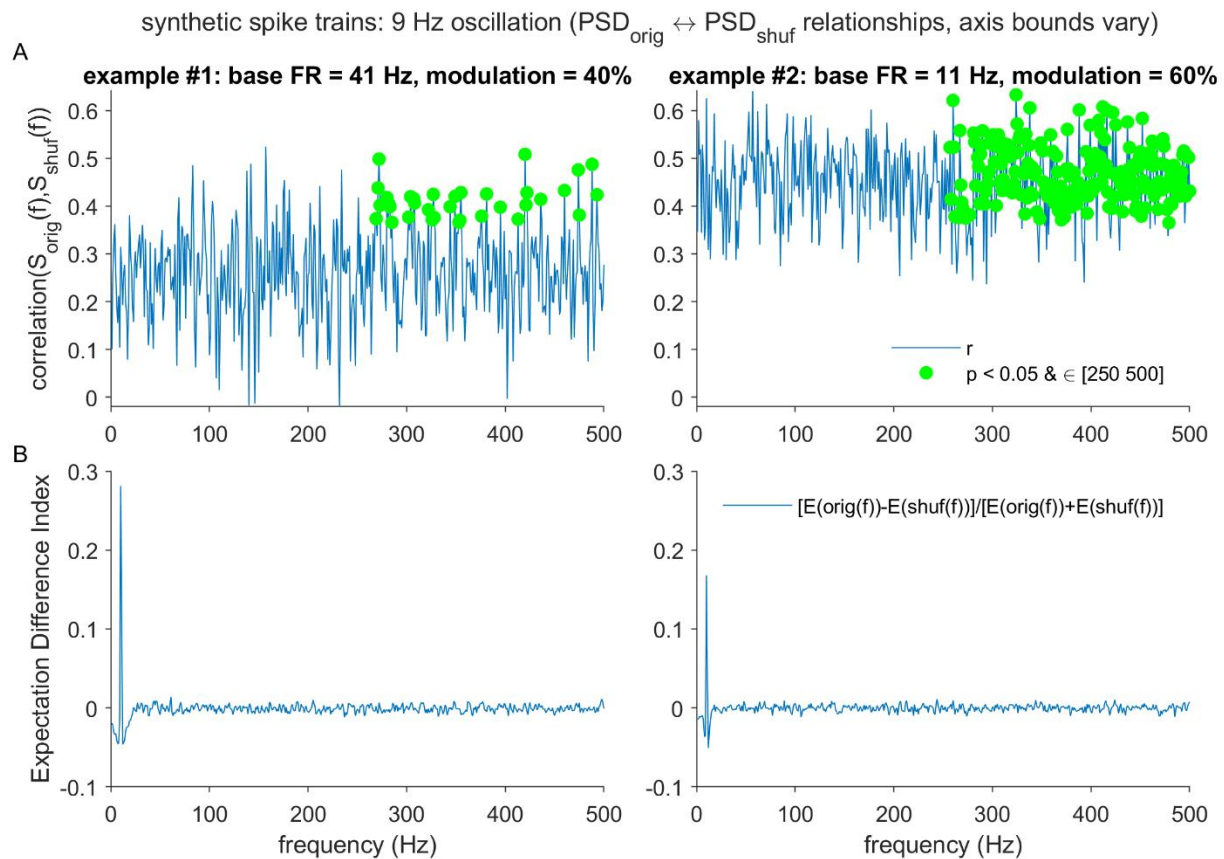
As described in Section 2.5, the final step of the shuffling method [7] entails the division of the original PSD by the mean of the shuffled spike trains' PSDs. One might ask why division, rather than subtraction, was used for this step. The authors argued that division produces the additional side benefit of uniform variance of the corrected spectrum across the frequencies. Such uniform variance would be desirable, since the mean and standard deviation of the corrected PSD over a subset of high frequencies is used to inform the construction of confidence intervals.

The authors offered a proof to support their claim that the division operation results in uniform variance. The reader is referred to Appendix B in [7] for details; for the present, note that the proof makes two testable assumptions. Adopting the notation from [7], let  $S_{orig}(f)$  and  $S_{shuf}(f)$  denote the values of the original PSD, and the mean of the shuffling-based PSDs, at frequency  $f$ . The first assumption states that  $\text{cov}(S_{orig}(f), S_{shuf}(f)) = 0$ ; that is, that the covariance of the original and mean shuffled PSDs, at each frequency, is zero. The second assumption (which implies a focus on the CI-generating high frequency range in particular) states that  $E(S_{orig}(f)) = E(S_{shuf}(f))$ ; that is, that the expected PSD for the original and mean shuffled PSDs are equal for all considered frequencies.

It is useful to ask whether these two assumptions are typically met in practice. Therefore, I carried out a series of assessments on a small synthetic dataset. For each of two sets of parameter assignments (matching those used to generate the Fig. 2-3 examples), I generated 100 new spike

trains. Using these two sets of spike trains, I generated estimates of  $\text{cov}(S_{orig}(f), S_{shuf}(f))$ , and  $[\text{E}(S_{orig}(f)) - \text{E}(S_{shuf}(f))]$ . Since the absolute values of these quantities are difficult to interpret, I plot normalized measures (the correlation coefficients,  $r$ , for  $S_{orig}(f)$  and  $S_{shuf}(f)$ , and an expectation difference index,  $\text{EDI} = [\text{E}(S_{orig}(f)) - \text{E}(S_{shuf}(f))] / [\text{E}(S_{orig}(f)) + \text{E}(S_{shuf}(f))]$ ), in Fig. A1.

Focusing on the 250-500 Hz control range specifically, the data appear to support only one of the two assumptions. For both simulation cases, several  $r$  values reliably exceeded 0 (fraction of  $r$  values significant at the corrected  $\alpha = .05$  level: 35/257 for the  $p_{base} = 41$  Hz,  $m = 0.4$  case; 221/257 for the  $p_{base} = 11$  Hz,  $m = 0.6$  case). Since a nonzero correlation implies nonzero covariance, the covariance assumption does not appear to hold for these cases, with the most notable violation present in the low FR example. In contrast, the EDI remained close to 0 in both cases, over the 250-500 Hz range.



**Figure A1. Simulation-based assessment of the assumptions of the proposed proof from Rivlin-Etzion et al. [11].** The authors aimed to support the claim that the ratio of the original ( $S_{orig}$ ) and shuffling-estimated ( $S_{shuf}$ ) PSDs yields a corrected PSD with uniform variance across frequencies, specifically within the control high-frequency range used for confidence interval (CI) construction. The proof assumed that the covariance of  $S_{orig}$  and  $S_{shuf}$  at each frequency  $f$  would be 0, and the expected values of  $S_{orig}$  and  $S_{shuf}$  at each  $f$  would be equal. Estimates of the relevant quantities were generated by creating 100 new spike trains for each of the two sets of simulation parameter assignments as were used to generate Figures 2-3. Normalized versions of these quantities are plotted to aid in value interpretation. (A) Pearson correlation coefficients between the 100 pairs of  $S_{orig}$  and  $S_{shuf}$  estimates, as a function of the oscillation frequency  $f$ . These  $r(f)$  estimates permit assessment of the assumption that  $cov(S_{orig}(f), S_{shuf}(f))$  is zero. Green dots indicate points that both fall within the 250-500 Hz band of interest, and for which  $r(f)$  significantly differs from 0 ( $p < .05$ , two-tailed, corrected for multiple comparisons). (B) Plots of the difference  $E(S_{orig}(f)) - E(S_{shuf}(f))$ , normalized by these expectations' sum. These plots permit assessment of the assumption that  $E(S_{orig}(f)) = E(S_{shuf}(f))$  over the 250-500 Hz frequency range.

The violation of the covariance assumption does raise doubt concerning the accuracy of the assertion that the division step will translate into uniform variances across the control range frequencies (especially for low firing rate cases). To the extent that the variances do differ, the reliability of the constructed confidence intervals is reduced. Since [7] did report that their simulation results remained largely unchanged when replacing the empirically-determined CIs with those generated by an analytical procedure [26], this reliability issue does not appear to pose a severe problem. However, a concerned researcher could elect to use the analytical approach (or verify result stability over repeated shuffling correction runs) if this issue is of concern.

## **Appendix A.2 Method Evaluation on an Alternative Simulation Framework**

Throughout the main text, I presented simulations generated through the framework introduced by [7]. In this framework, the spike rate oscillation arises from a pure (i.e., noiseless) sine function, and the recovery period adopts an exponential form. One may raise the question of how the residuals and shuffling methods might perform over synthetic data generated by more complex, and potentially more biologically-realistic frameworks. This section of the Appendix describes the methods and results for a limited trial of an alternative simulation framework.

In assembling this alternative framework, I target two spiking features. First, I sought to replace the deterministic sine term from [7] with a stochastic oscillation. Such oscillations may remain anchored on a sine wave of a specific center frequency, but exhibit noise-driven departures from this precise form, resulting in temporal fluctuations of amplitude and phase [76]. Second, I aimed to replace the assumed exponential recovery period with the trajectory that emerges from a standard biophysical model of neural spiking (i.e., one which recreates the channel dynamics that gives rise to a cell's intrinsic RP). I detail the specific framework components that address these two goals in the following Methods section.

### **Appendix A.2.1 Simulation Methods**

As an overview, the alternative framework drew upon three existing models, each of which provided a useful feature. To generate the stochastic oscillations, I adopted the system of stochastic differential equations (SDEs) presented by [76]. These authors used their generated oscillations to drive a “perfect integrate-and-fire” (PIF) unit (see below), which is not a biophysical

model (and lacks a recovery period component). To model the driven neuron (which generated the spikes to which spectral analysis were applied), I used the Hodgkin-Huxley equations [77], which are broadly recognized as the foundational biophysical model of neuronal spiking. Finally, adopting elements from [11], I connected the oscillatory drive to the target neuron by (a) using the stochastic oscillation to drive the spike rates of 200 afferent (i.e., upstream, “source”) neurons, which in turn acted through a model synapse to drive the target’s spiking.

The full model of [76] consists of four equations. The first equation instantiates the PIF neuron (and its right-hand side provided the drive that I used to modulate the 200 source neurons). The PIF equation models the ongoing change in the neuron’s voltage,  $\dot{v}(t)$ , as the sum of a base firing rate ( $\mu$ ) and additional driving forces ( $x(t)$ ,  $z(t)$ ):

$$\dot{v}(t) = \mu + x(t) + z(t) \tag{9}$$

The neuron is said to spike when  $v(t)$  exceeds a voltage threshold ( $v_T$ ), with  $v(t)$  reset to 0 on the subsequent time step. Of the various discrepancies between the PIF and more biologically realistic models, it is useful here to note that the PIF lacks an explicit recovery period.

The first external drive term,  $x(t)$ , represents the stochastic oscillation, which is in turn generated by the following system of equations (where  $\xi(t)$  represents scaled, normally distributed noise):

$$\dot{x}(t) = y(t) \tag{10}$$

$$\dot{y}(t) = -\gamma y(t) - w_0^2 x(t) + \sqrt{(2D)}\xi(t) \tag{11}$$

In exploring their model, the authors described indirect manipulation of the parameters of Eq. (11) through direct variation of three unshown parameters:

1.  $w$ : the “frequency ratio”, which may be used to define the central frequency of the oscillation, relative to the mean spike rate  $\mu/v_T$
2.  $Q$ : the “quality factor”, which determines the oscillation’s bandwidth and coherence
3.  $\widehat{\sigma_x^2}$ : the variance of the stochastic oscillation

Given settings for the three above parameters, the three parameters present in (11) are set as  $\gamma = 2\pi w\mu/Qv_T$ ,  $w_0^2 = (2\pi w\mu/v_T)^2(1 + 1/4Q^2)$ , and  $D = \gamma\omega_0^2\mu^2\widehat{\sigma_x^2}$ .

The second external drive,  $u(t)$ , represents broadband noise, as generated by the following Ornstein-Uhlenbeck (OU) process (where  $\xi_2(t)$  represents an independent sample of scaled Gaussian noise):

$$\dot{z}(t) = -z(t)/\tau + \sqrt{(2D_z/\tau^2)}\xi_2(t) \tag{12}$$

Although this broadband noise term is not necessary for the goal of evaluating the methods’ detection of stochastic oscillations, I include it (with a low  $\widehat{\sigma_z^2}$  setting) to further contribute to the realism of the model.

In a manner similar to Equation (11), the authors describe the manipulation of the free parameters for Equation (12) indirectly, through setting of the temporal correlation-governing parameter  $\hat{\tau}$  (such that  $\tau = \hat{\tau}v_T/\mu$ ) and the OU process variance  $\widehat{\sigma_z^2}$  (such that  $D_z = \mu v_T \widehat{\sigma_z^2} \hat{\tau}$ ).

It is important to note that the dynamics of  $x(t)$ ,  $y(t)$ , and  $z(t)$  unfold independently of the status of  $v(t)$ ; the occurrence of a spike has no impact on these variables.

For each new simulation run that I executed, I first generated  $T = 30 \times 1024$  ms of synthetic data using the [76] model, with the results of the Eq. (9) sum stored in a variable  $s(t)$ , rather than being used to drive a PIF unit. Simulations were carried out in XPP [78] using the Euler–Maruyama method [79], and a step size of  $\Delta t = 1e^{-4}$  s. Although I varied some parameters over runs, others remained fixed, such that  $v_T = 1$ ,  $\widehat{\sigma_z^2} = 0.0025$ , and  $\hat{\tau} = 1$ . For the results reported here, I tested a center oscillation frequency of 20 Hz, implying that  $w = 20/\mu$ .

The  $s(t)$  output was conceptualized as an oscillatory rate function that was shared by all source units. I generated 200 novel instances of the inhomogeneous Poisson process defined by this function, resulting in a set of source neuron spike trains. This means of creating the source units' firing is similar to the procedure described in [80] (although I omitted these authors' inclusion of additional rate function components that were unique to individual units or subgroups of units). As in [80], no recovery period was introduced for these source units, as they are not intended to represent realistic behaviors of individual neurons so much as the noisy oscillatory input that emerges from an entire source population.

In generating the source spike trains, I converted the PIF input drives represented in  $s(t)$  to spike probabilities,  $p_{spk}(t)$ , for each time bin. For some simulation outputs, all values of the product  $s(t)\Delta t$  fell in  $[0,1]$ , and this product was used to set the corresponding  $p_{spk}(t)$  values. However,  $s(t)\Delta t$  values could fall outside  $[0,1]$  (in practice, this always involved the occurrence of negative



$s(t)$  values). To perform the conversion of  $s(t)$  inputs that could not be directly mapped to spike probabilities, I set  $p_{spk}(t) = (s(t)-\bar{s})(\bar{s}/(\bar{s} - \min(s)) + \bar{s}$ , so that the mean  $\bar{s}$  was preserved, but the range was compressed to the extent needed to ensure all  $p_{spk}(t) \geq 0$ .

Let  $t_k^j$  represent the timestamp of the  $k$ th spike produced by the  $j$ th source neuron. As in [80], these timestamps were understood to correspond to the release of neurotransmitter-containing vesicles into the synapse between each source and target neuron. While referring the reader to general neuroscientific texts for more details on synaptic transmission (e.g., [1]), it is sufficient to note that vesicle release is assumed to impact the postsynaptic target neuron by altering the conductivity of its membrane to specific forms of ion flow. The above said, the next step in simulating modulatory drive of the target unit entails conversion of the spike arrival timestamps to the resulting post-synaptic conductance.

Following from [80] (and see also [16, 81], for general references) the postsynaptic conductance produced by the  $j$ th source neuron is modeled as

$$g_j(t) = \bar{w} \sum_k \alpha(t - t_k^j) \tag{13}$$

where  $\bar{w}$  is a fixed weight (varied across settings to simulate target neurons of varying firing rates; see descriptions in section A.2.2). The function  $\alpha(t - t_k)$  models the rise and decay of the conductance attributable to spike  $k$ :

$$\alpha(t) = \frac{J}{\tau_2 - \tau_1} \Theta(t)(e^{-t/\tau_2} - e^{-t/\tau_1}) \tag{14}$$

Here,  $\tau_1 = 4$  ms,  $\tau_2 = 1$  ms,  $\Theta(t)$  is the Heaviside step function, and  $J = 10^{-4}$ .

The  $g_j(t)$  values are then summed over the source units, to obtain the total postsynaptic conductance contribution  $G(t)$ :

$$G(t) = \sum_{j=1}^{200} g_j(t) \quad (15)$$

The postsynaptic conductance,  $G(t)$ , results in current flow into the target cell. Again following conventional methods [16, 80], this synaptic current,  $I_{syn}$ , was modeled as

$$I_{syn} = g(t)(V - V_{rev}) \quad (16)$$

where  $V$  is the current membrane potential of the postsynaptic, target cell, and  $V_{rev}$  (set here to 0) is the reversal potential (the voltage at which the direction of current flow reverses; see [16, 81] for details).

This  $I_{syn}$  term is incorporated into the Hodgkin-Huxley (HH) model of the target neuron. Full discussion of the Hodgkin-Huxley equations is beyond the scope of thesis; see [77] for the original presentation. For the present purposes, it is sufficient to be aware that the HH equations can be used to model (1) the dynamics of the membrane potential,  $V$ , as observed before, during, and after a spike (including during the absolute and relative refractory periods), and (2) the impact of externally applied current (such as  $I_{syn}$ ) on  $V$ . The equation for  $V$  takes on the form

$$C_m V' = -I_{Na} - I_K - I_L - I_{syn} \quad (17)$$

where  $C_m$  is the membrane capacitance, and the  $I$  terms refer to the  $\text{Na}^+$ ,  $\text{K}^+$ , leak, and synaptic currents [78].

I implemented the HH model in XPP (step size  $\Delta t = 1e^{-4}$ , Gear method), with  $I_{syn}$  set according to the output from (16), and all other adjustable parameters set according to values from [78]. The resulting voltage time series,  $V(t)$ , was exported to Matlab for further processing. The timestamps of action potentials were first identified by locating local maxima in  $V(t)$  that exceed 20 mV. These timestamps were then used to prepare a binary delta vector representation of the spike train (downsampled to the conventional 1 kHz resolution), as described in Chapter 3. The delta vector was submitted to spectral analysis through the uncorrected, shuffling, and residuals methods.

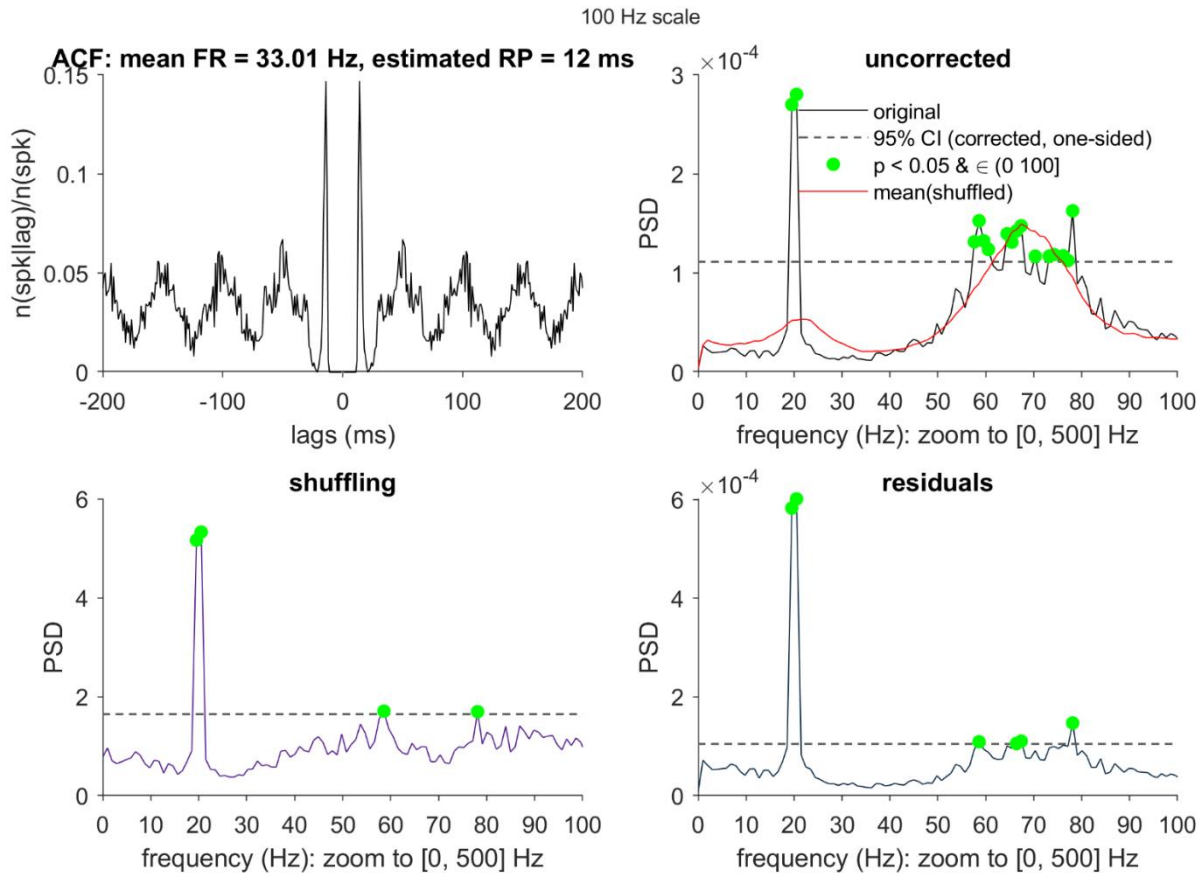
### Appendix A.2.2 Simulation Results

Below, I report on the performance of the PSD correction methods on example cases. In each case, a “hit” occurs if the method enables detection of an oscillation near 20 Hz. For generation of the stochastic oscillations, I used parameter settings that were of similar scale as those reported in the original model from [76], and base any remarks on the “high” or “low” status of any settings on comments in their report.

For the first example, I generated a stochastic oscillatory drive for which the conditions for detection were favorable. The oscillatory drive was generated with a relatively high base FR to center frequency ratio ( $1/w = 2$ ), a high quality factor ( $Q = 40$ ), a typical variance setting ( $\widehat{\sigma}_x^2 = .04$ ), and a synaptic weight that favored a moderate firing rate in the target neuron ( $\bar{w} = 0.2$ ). The PSDs for the resulting spike train are shown in Figure A2.

The results prompt four observations: (1) Based on the autocorrelation function, estimation of the RP duration appears to remain reasonable, even though I have replaced the assumed exponential recovery with the form that emerges from a biophysically-grounded model. (2) As

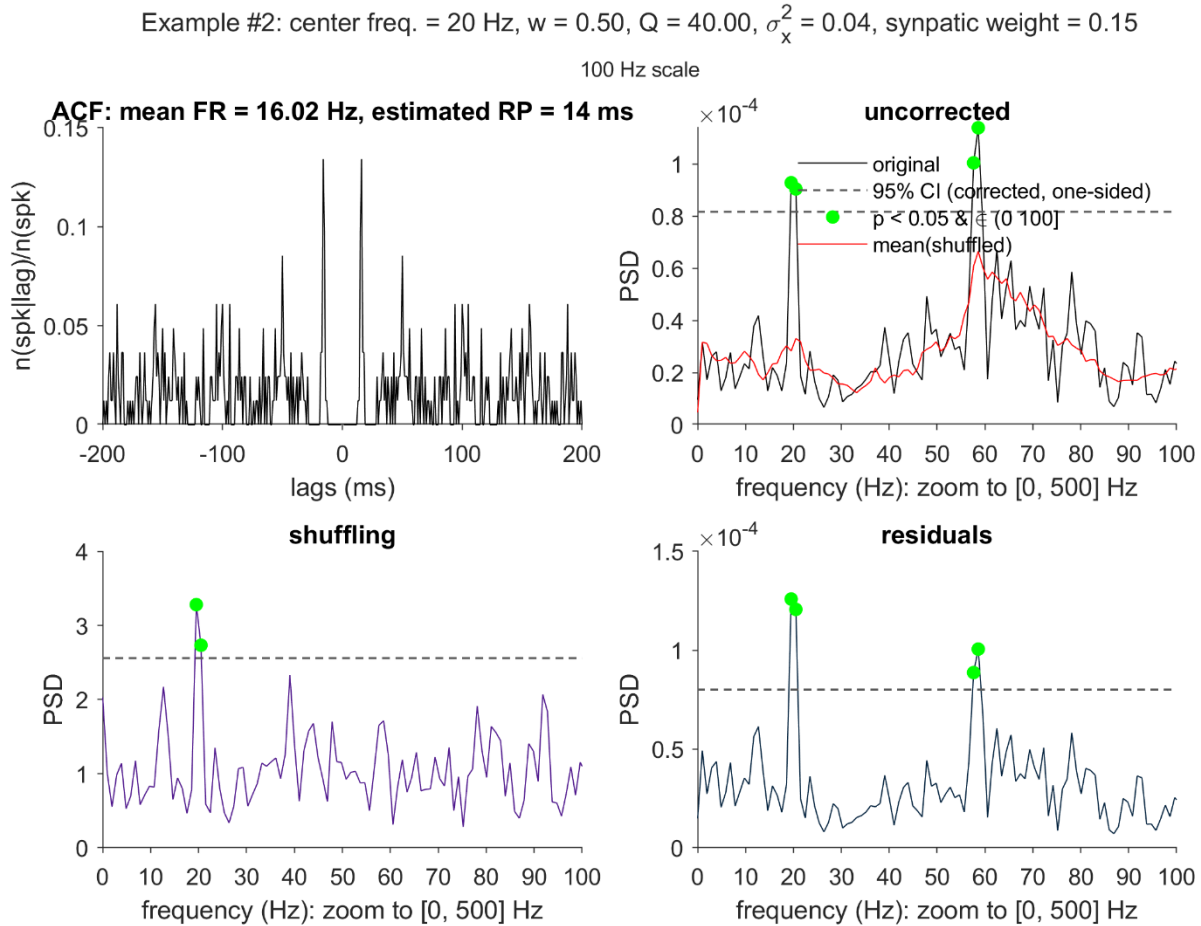
Example #1: center freq. = 20 Hz,  $w = 0.50$ ,  $Q = 40.00$ ,  $\sigma_x^2 = 0.04$ , synaptic weight = 0.20



**Figure A2. Output of the three PSD methods (uncorrected, shuffling, and residuals) on the first example spike train generated under the alternative simulation framework.** The example was created with parameter settings intended to produce a readily detectable oscillation (i.e., high quality, with moderate FRs for both the driving oscillation and target neuron). See the text for definitions of the manipulated simulation parameters.

anticipated given the target unit's FR and the high oscillation quality, all PSD methods (uncorrected, shuffling, residuals) detected the 20 Hz oscillation. The addition of noise to the oscillatory dynamics was not an impediment. (3) In the uncorrected PSD, very strong distortion is apparent over the high frequencies, likely owing to the broad absolute component of the RP generated by the HH model (evident as the 0 values present in the center of the ACF plot). (4) Over this high-frequency region, both the shuffling and residuals methods return false alarm reports of significant oscillation. This pattern is similar to that described for the high FR, strong modulation cases reported in Chapter 4.

For Example #2, I asked how the results pattern might change when the synaptic weight  $\bar{w}$  is reduced to 0.15 (**Error! Reference source not found.**). This change implies that the oscillatory drive is expressed through a more slowly spiking target unit.



**Figure A3. Output of the three PSD methods on the second example spike train generated under the alternative simulation framework.** Relative to Example #1, this spike train was produced using a lower synaptic weight setting ( $\bar{w} = 0.15$ ), translating into a lower firing rate for the unit driven by the upstream oscillation.

In these results, three observations stand out. First, in spite of the reduced firing rate, the strength of the oscillatory drives implies that all methods continue to detect the 20 Hz oscillation.

Second, the shuffling method does show some evidence of removing some of the PSD mass associated with the oscillation (see red line in the uncorrected PSD panel). However, in this

particular case, this erosion of the peak did not prevent detection of the oscillation in the shuffling-corrected PSD.

Finally, the false alarms over the high-frequency points – prevalent in all PSDs in Figure A2 – only remained present in the uncorrected and residuals PSDs in Fig. A3. This finding differs from the residuals false-alarm trends observed in Chapter 4, as the Fig. A3 false alarms were generated in the presence of a low firing rate, and for a moderate oscillation frequency (while the previous observations occurred in the presence of high firing rates, and predominantly over the low oscillation frequencies). Future simulation-based evaluations would be useful to determine if the appearance of the false alarms in the residuals, but not the shuffling output, is a reliably replicable result in this simulation setting. If this the case, then one speculative diagnosis for the residuals method’s difficulty concerns the “perfect predictor” problem [82]. Perfect prediction arises as a potential issue in the present context due to the extended period of absolute refractoriness mentioned above. For each post-spike time bin that is both  $\leq \hat{n}_r$  (i.e., within the modeled RP) and corresponds to zero spike probability, the corresponding GLM regressor would be considered a “perfect” predictor, referring to the fact that an infinite range of negative weight assignments for this term would result in perfect prediction of the zero spike rate. This situation can create vulnerability to model instability.

In light of the above, it would be useful to further explore whether the residuals method is reliably challenged by RPs with prolonged intervals of zero spike probability, and if so, add improvements to address this challenge [82]. Although one might question how frequently the HH model’s prolonged interval of exactly zero spiking would be observed in real neurons (see the Chapter 2 comments regarding the typical brevity of empirically measured absolute

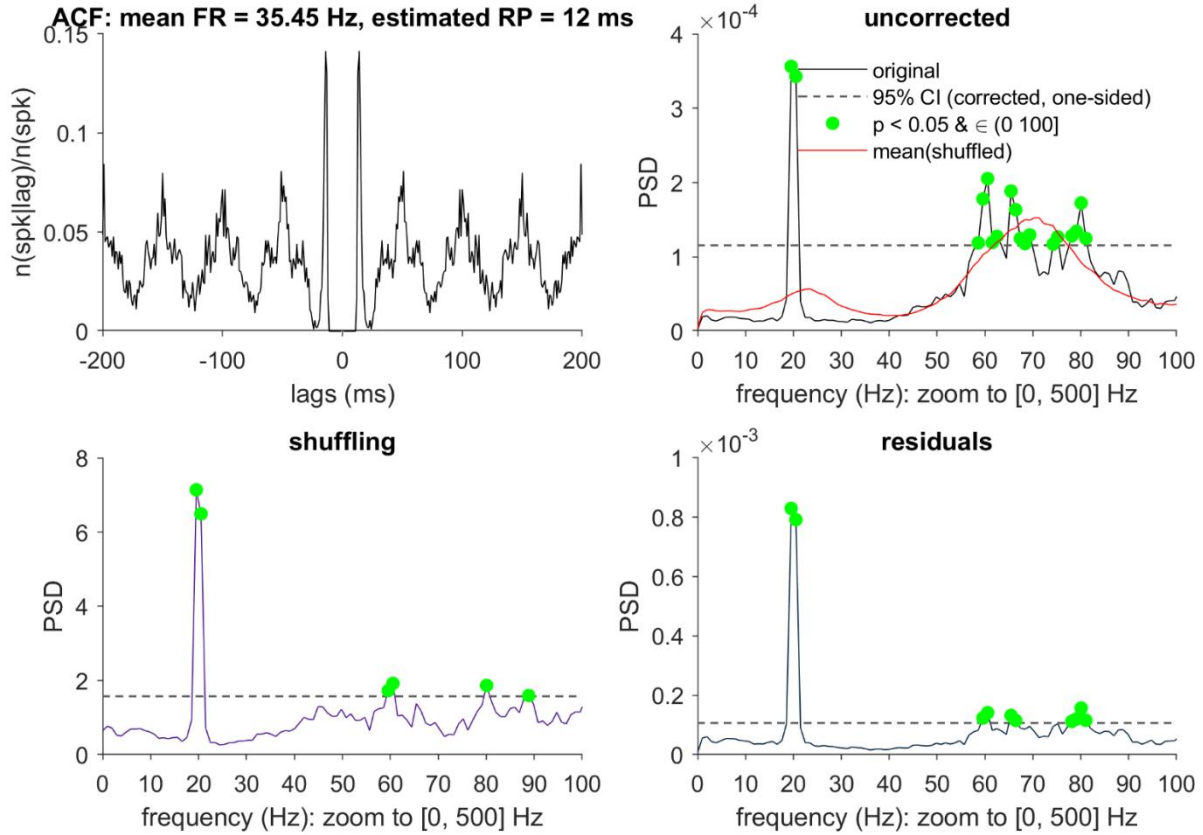
refractoriness), accommodation of this pattern by the residuals framework might improve its applicability to other domains, in which long pauses of zero event arrivals might be common.

For Example #3, I again applied the parameter settings from Example #1, with the exception of the  $\hat{\sigma}_x^2$  setting, which was increased from 0.04 to 0.16 (that is,  $\hat{\sigma}_x$  was doubled). This change introduces additional noise into the stochastic oscillatory drive. Comparing the results (Figure A4.) with those of Figure A2, it appears that this parameter increase did not meaningfully change the major characteristics of the uncorrected or corrected PSDs.

Overall, these case studies support a common conclusion: When transitioning from our original framework (which operated at a high level of abstraction, capturing only the core features of the process of interest) to a more complex approach (which attempted to incorporate more biologically-relevant details), we saw that our major observations regarding the residuals and shuffling methods remained largely consistent.

Example #3: central frequency = 20 Hz,  $w = 0.50$ ,  $Q = 40.00$ ,  $\sigma_x^2 = 0.16$

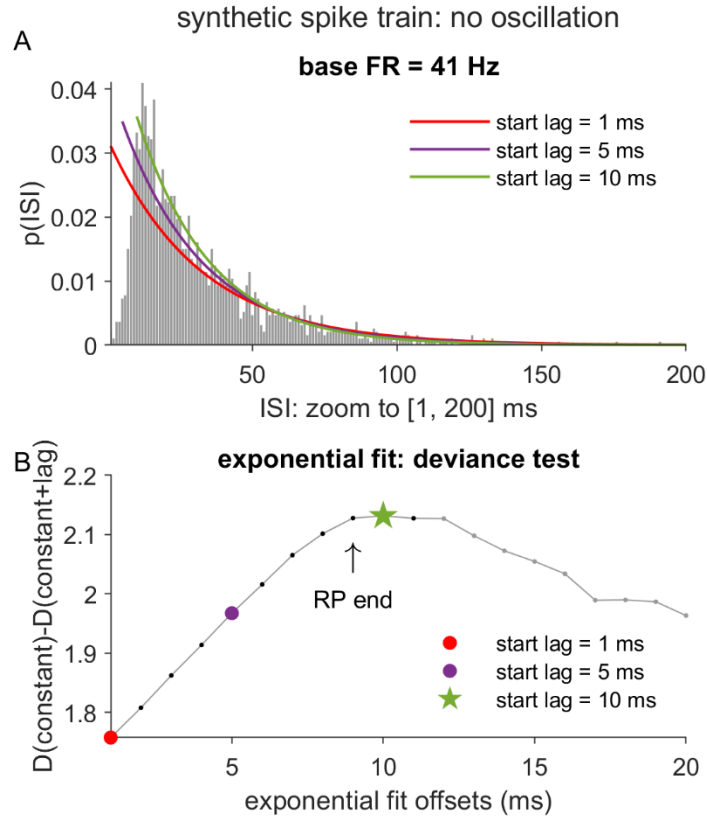
100 Hz scale



**Figure A4. Output of the three PSD methods on the third example spike train generated under the alternative simulation framework.** Relative to Example #1, this spike train was produced under conditions of increased noise in the underlying oscillation trend ( $\hat{\sigma}_x^2 = 0.16$ ).

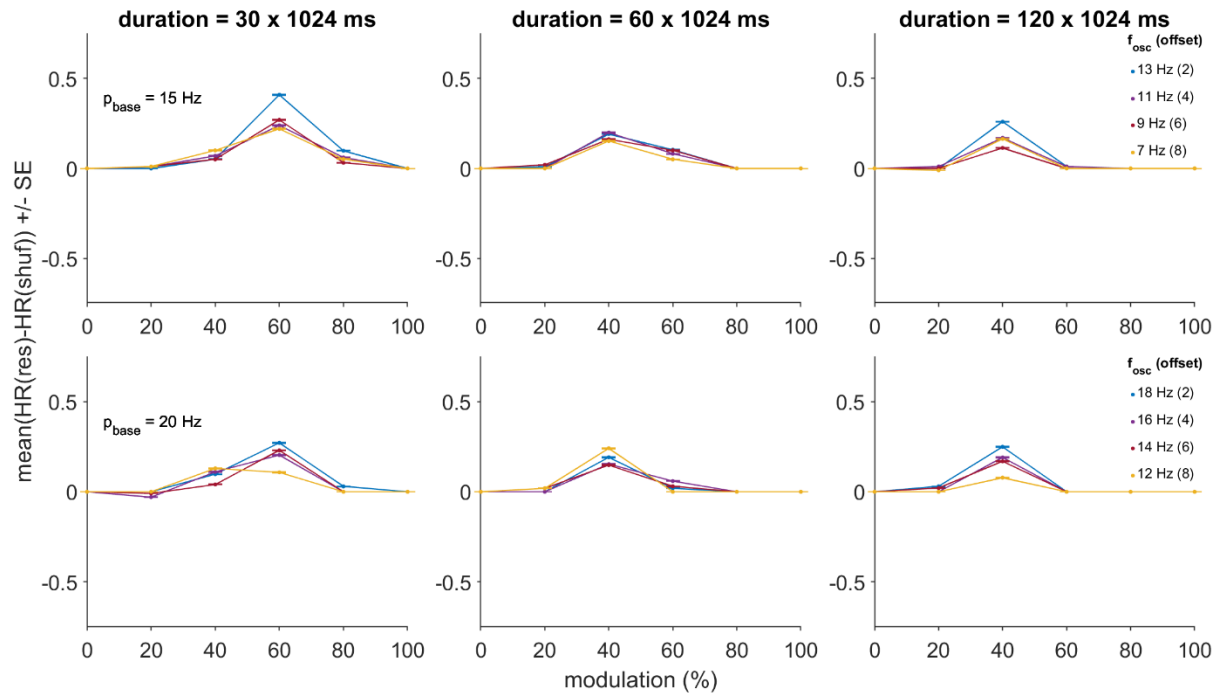


## Appendix A.3 Supporting Figures



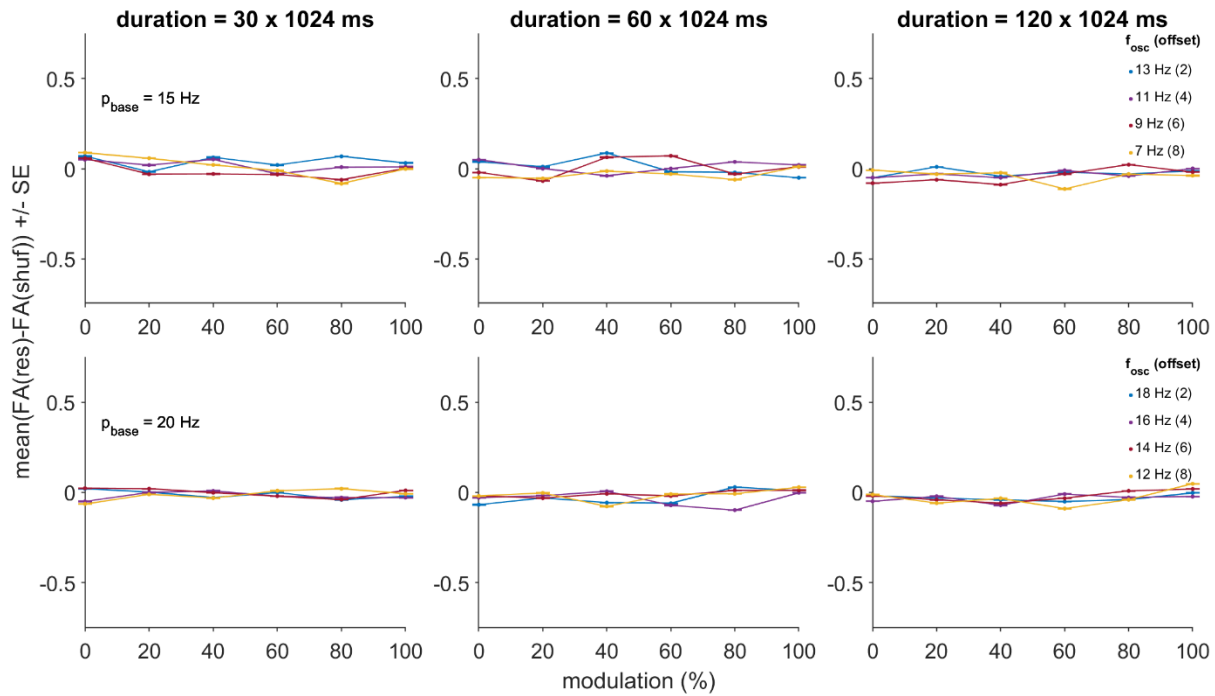
**Figure A5. Estimation of recovery period duration: No oscillation example.** (A) Illustration of the procedure for obtaining an estimate of the RP duration ( $\hat{n}_r$ ), as applied to a synthetic spike train with no oscillation (modulation strength  $m = 0$ ). A series of right-shifted exponential curves are fit to the ISI distribution, left-anchored to starting positions advanced in 1 ms steps (with 3 sample iterations highlighted in the figure). (B) Plot of the deviance difference statistic,  $\Delta D$ , as a function of the first 20 starting positions of the exponential fits.  $D(\text{constant})$ ,  $D(\text{constant}+\text{lag}) =$  deviance measures for the intercept-only and intercept+exponential curve models, respectively.  $\Delta D$  tracks the goodness of fit contributed by the exponential curve. The  $\hat{n}_r$  estimate is set equal to the post-spike lag immediately preceding the first local maximum in the  $\Delta D$  plot.

synthetic data: hit rate differences (residuals - shuffling),  $p_{\text{base}}$  varied directly (9 ms RP,  $k = 0.7$ )



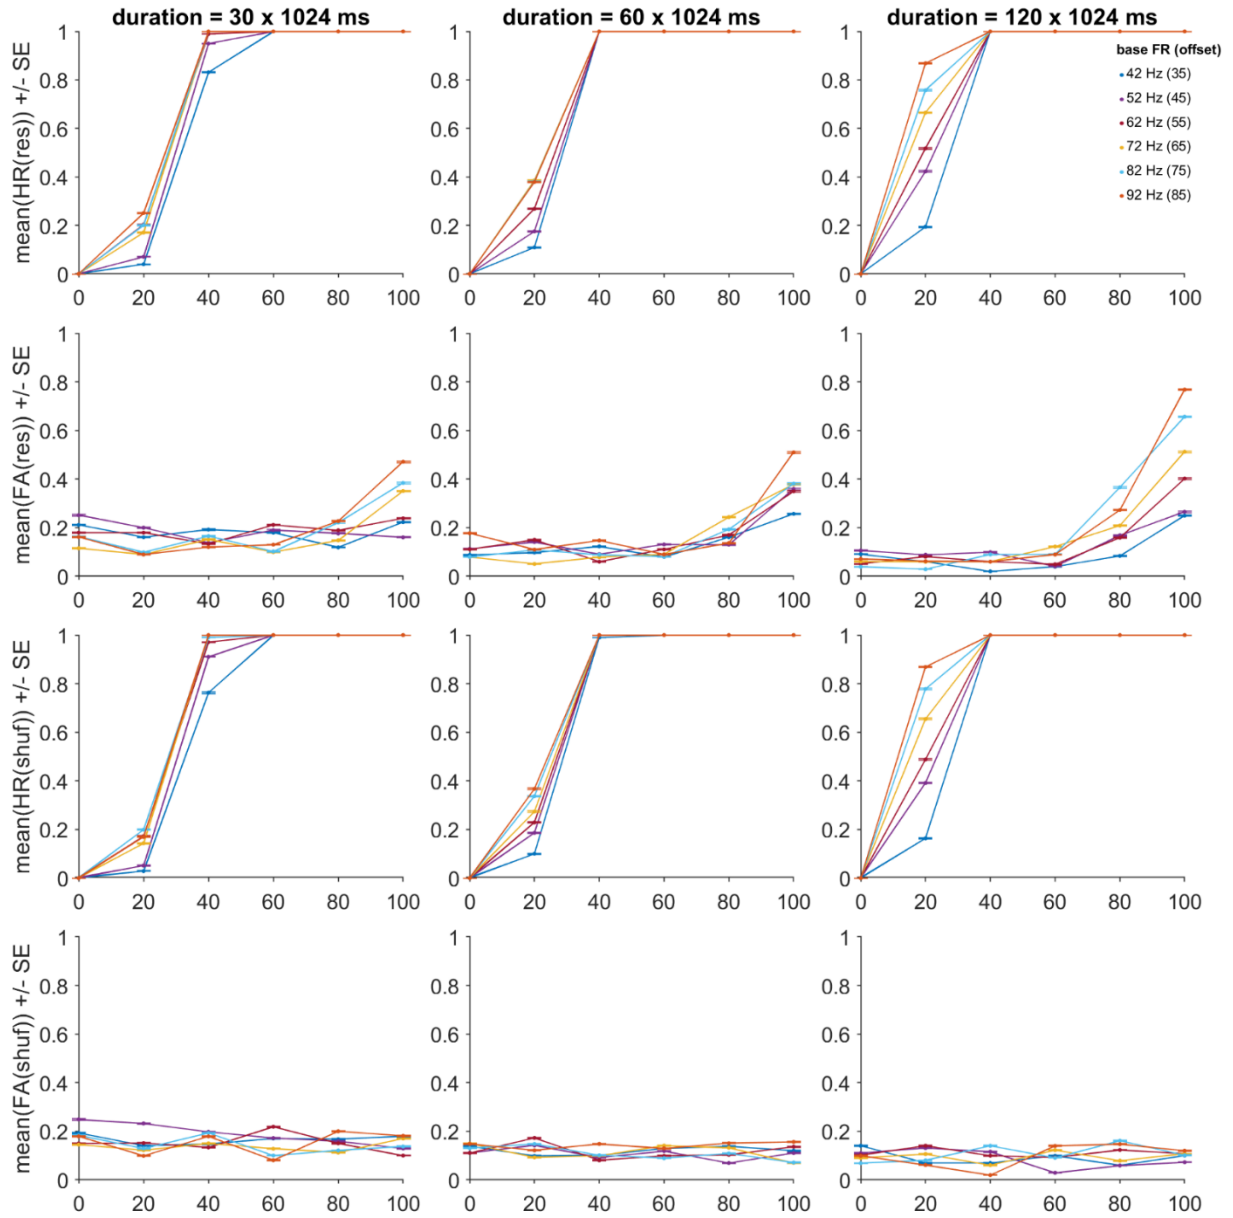
**Figure A6. Residuals - shuffling difference in hit rates over the varied parameters of the dataset formed with direct base firing rate manipulation.** Means and standard errors reflect summaries over the hit rates computed for each of 1000 subsamples of a dataset in which the pbase (i.e., base FR) and pbase\_offset parameters were varied directly (as opposed to  $p_{\text{base\_offset}}$  and  $f_{\text{osc}}$ ). Plots depicts all 144 unique combinations of pbase (rows), simulation duration ( $T$ , columns), oscillation modulation strength ( $m$ , x axes), and oscillation frequency ( $f_{\text{osc}}$ , lines). All other abbreviations follow from those described for Figure 5. See Chapter 3 for a description of the subsampling procedure for this dataset.

synthetic data: false alarm differences (residuals - shuffling),  $p_{base}$  varied directly (9 ms RP,  $k = 0.7$ )



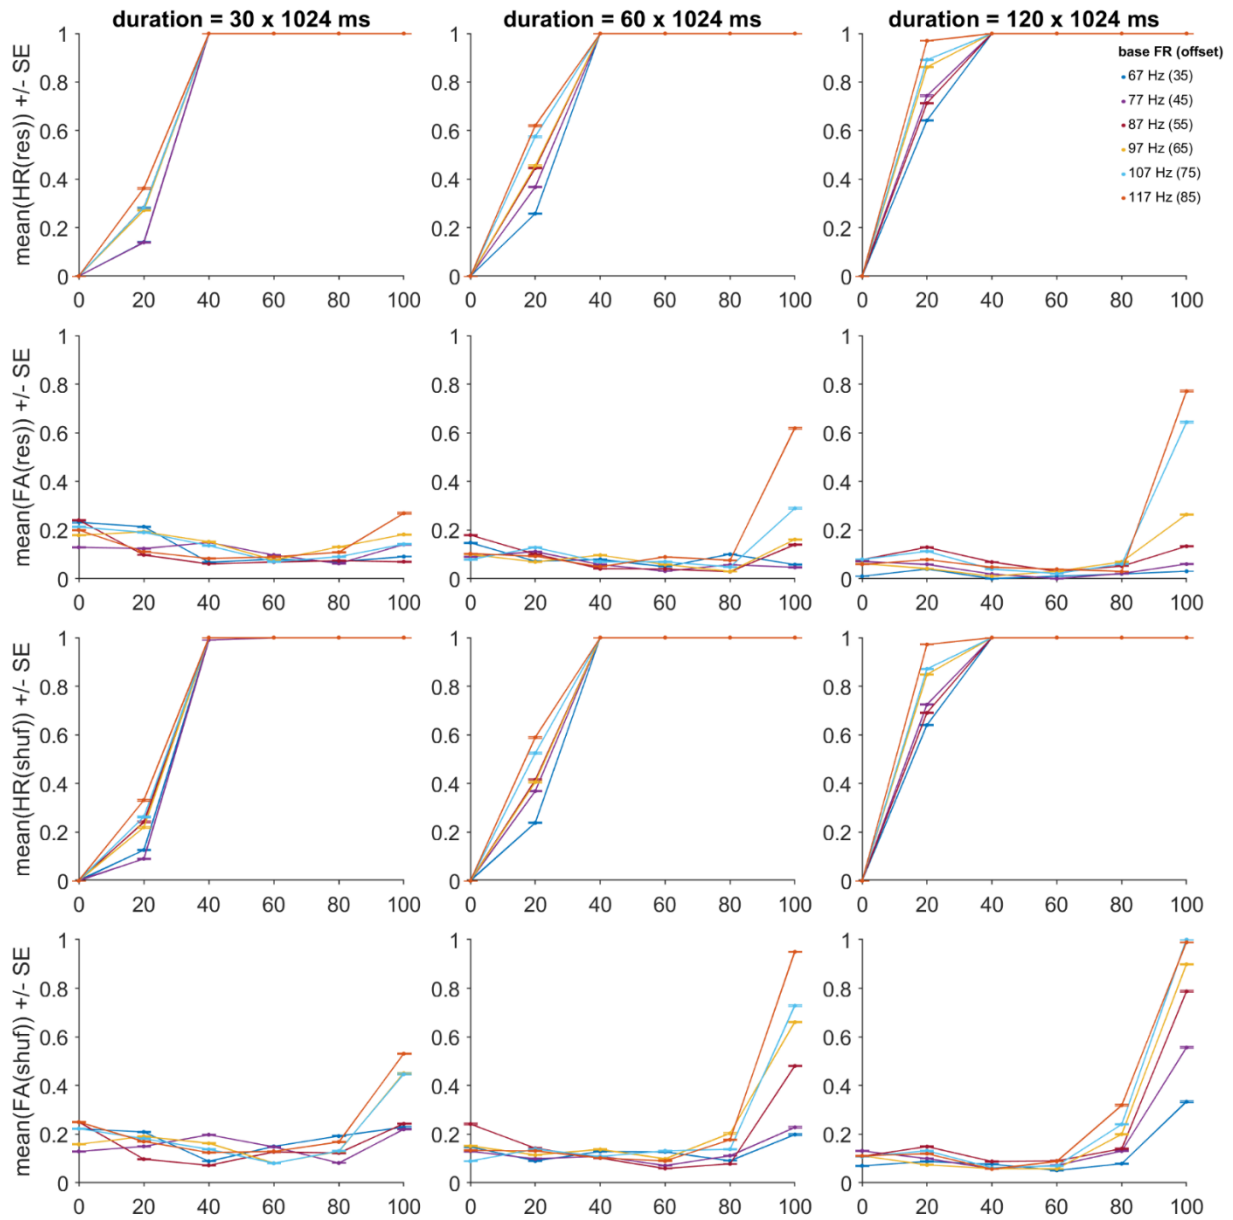
**Figure A7. Residuals – shuffling difference in false alarm rates over the varied parameters of the dataset formed with direct base firing rate manipulation.** Means and standard errors reflect summaries over the false alarm rates computed for each of 1000 subsamples of a dataset in which the  $p_{base}$  and  $p_{base\_offset}$  parameters were varied directly. All other abbreviations, the plotting conventions, and the subsampling procedure follow from those described for **Figure 6**.

synthetic data: method hit rates and false alarms for  $f_{osc} = 7$  Hz (high FRs, 9 ms RP,  $k = 0.7$ )



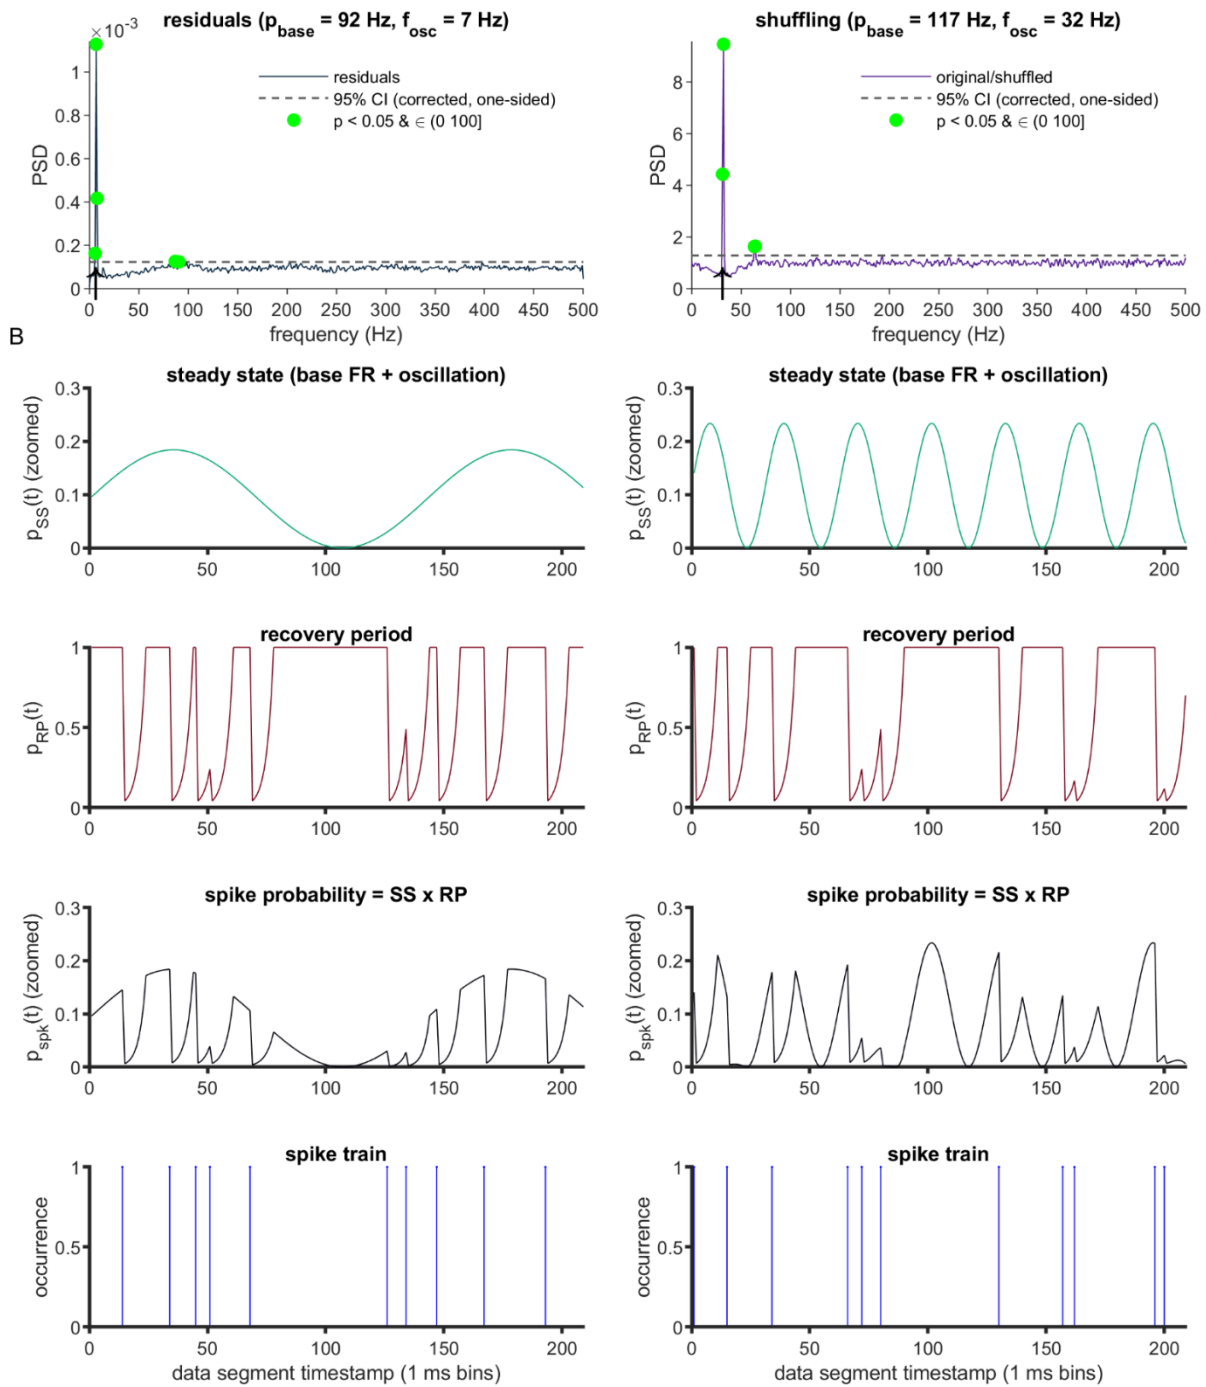
**Figure A8. Residuals and shuffling hit and false alarm rates for the dataset of high firing rate spike trains, with the oscillation frequency ( $f_{osc}$ ) fixed at 7 Hz.** Means and standard errors reflect summaries over the hit and false alarm rates computed for each of 1000 subsamples of the original, high FR dataset. Plots depict all 108 unique combinations of simulation duration ( $T$ , columns), oscillation modulation strength ( $m$ , x axes) and base FR - oscillation frequency offset ( $p_{base\_offset}$ , lines) for the cases in which  $f_{osc} = 7$  Hz. Abbreviations follow from those described for Figure 5-Figure 7.

synthetic data: method hit rates and false alarms for  $f_{osc} = 32$  Hz (high FRs, 9 ms RP,  $k = 0.7$ )



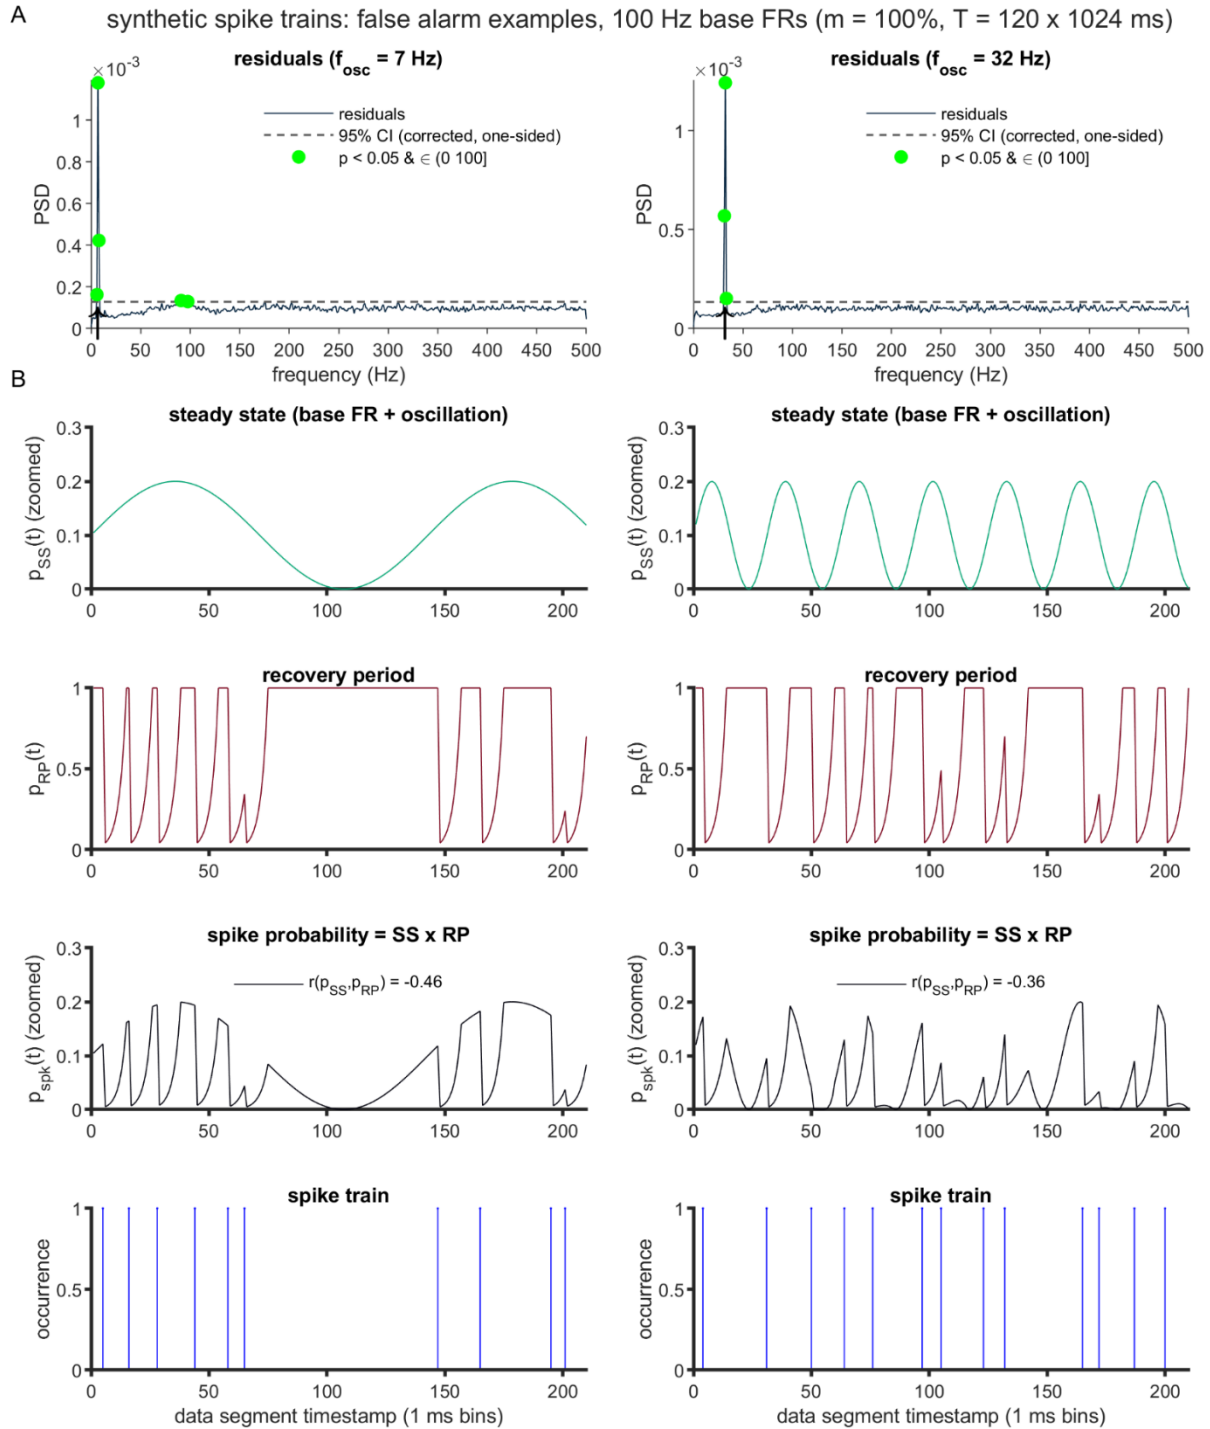
**Figure A9. Residuals and shuffling hit and false alarm rates for the dataset of high firing rate spike trains, with the oscillation frequency ( $f_{osc}$ ) fixed at 32 Hz.** Abbreviations and plotting conventions follow from those described for Figure 5-  
Figure 7.

A synthetic spike trains: false alarm examples, high base FRs ( $m = 100\%$ ,  $T = 120 \times 1024$  ms)



**Figure A10. Example power spectra and spike train segments for high firing rate cases that produced pronounced false alarms following residuals or shuffling correction.** Two synthetic spike trains were generated to illustrate conditions that yielded especially high false alarm rates under residuals and shuffling correction, respectively. Both spike trains utilized the maximal modulation strength ( $m = 1.0$ ), duration ( $T = 120 \times 1024$  ms), and base FR - oscillation frequency offset ( $p_{base\_offset} = 85$  Hz) settings from the original high FR dataset, and the default RP parameters (duration  $n_r = 9$  ms, steepness  $k = 0.7$ ). Oscillation frequency was

set at either the lowest setting from the full high FR dataset (7 Hz, to illustrate residuals FAs) or the highest setting (32 Hz, to illustrate shuffling FAs). (A) Corrected power spectral density (PSD) functions generated by the residuals method (left) and shuffling method (right). Statistical testing and plotting conventions follow from those described for Figure 2-Figure 3. (B) Illustration of the initial 209 ms of the synthetic spike trains (fourth row) from which the PSDs in (A) were computed, and the components of the rate functions that governed their generation (first-third rows; see Chapter 3, Eq. 1 for full details). Term definitions:  $pss(t)$  = steady state spiking probability,  $pRP(t)$  = recovery period spiking probability,  $pspk(t)$  = spike probability, reflecting the product of the steady state and recovery period components.





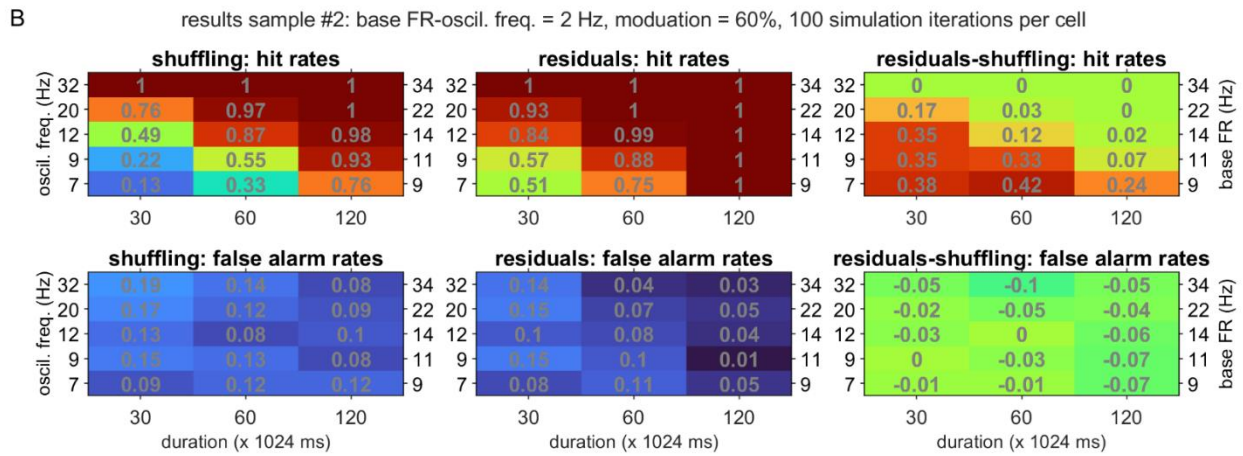
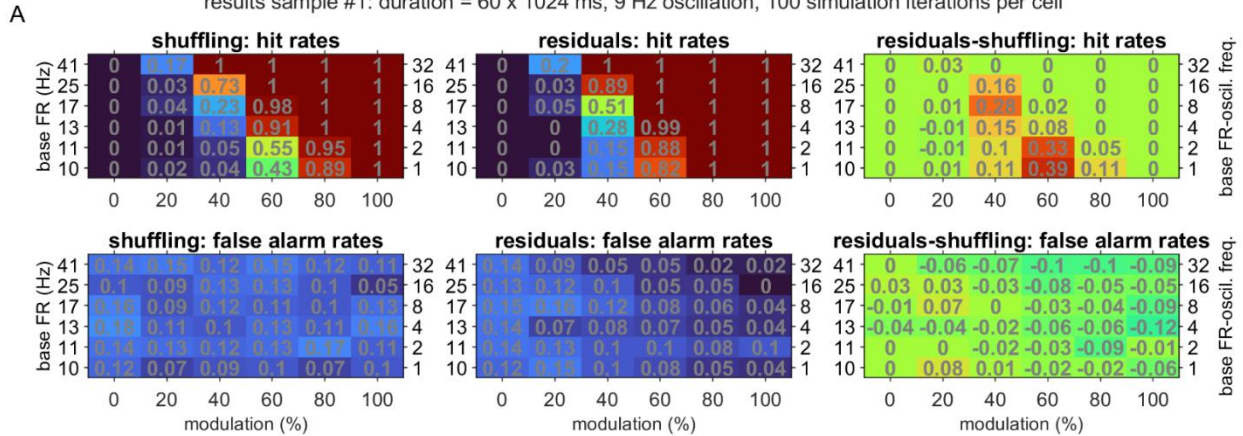
**Figure A11. Example residuals-corrected power spectra, and corresponding spike train samples, for two levels of oscillation frequency ( $f_{osc}$ ).** Two synthetic spike trains were generated using the same baseline FR (100 Hz), strong oscillatory modulation ( $m = 1.0$ ), duration ( $T = 120 \times 1024$  ms), and RP parameters (duration  $n_r = 9$  ms, steepness  $k = 0.7$ ), with strong oscillatory modulation ( $m = 1.0$ ) and high baseline firing rates (100 Hz). (A) Corrected power spectral density (PSD) functions generated by the residuals method for spike trains generated with  $f_{osc} = 7$  Hz (left) or 32 Hz (right). Statistical testing and plotting conventions following



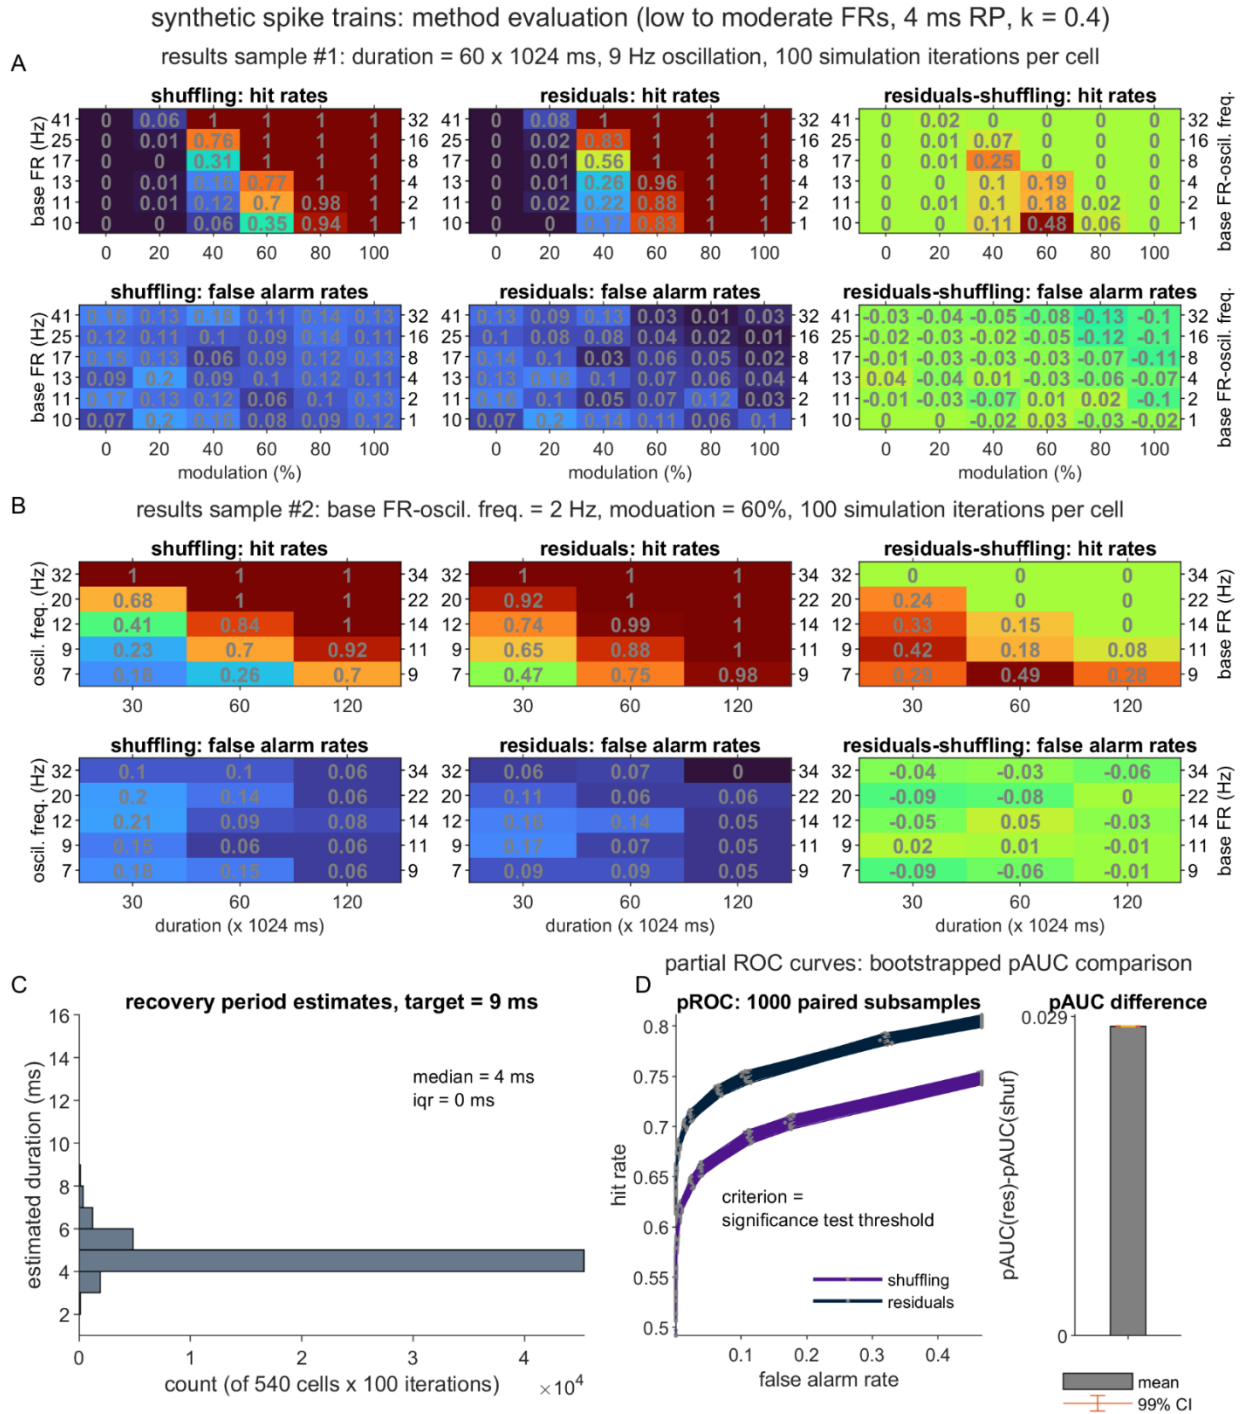
those described for Figure 2-**Figure 3**. (B) Illustration of the initial 209 ms of the synthetic spike trains (fourth row) from which the PSDs in (A) were computed, and the components of the rate functions that governed their generation (first-third rows; see **Error! Reference source not found.** for term definitions). The  $r$  values reported in the third row report the correlation between the recovery period contribution to the rate function ( $p_{RP}$ ) and the combined base FR and oscillatory contributions ( $p_{SS}$ , where SS = “steady state”).

synthetic spike trains: method evaluation (low to moderate FRs, 3 ms RP,  $k = 0.0$ )

results sample #1: duration = 60 x 1024 ms, 9 Hz oscillation, 100 simulation iterations per cell

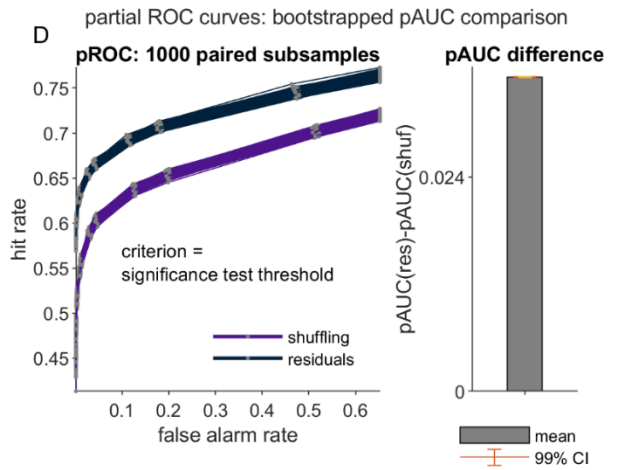


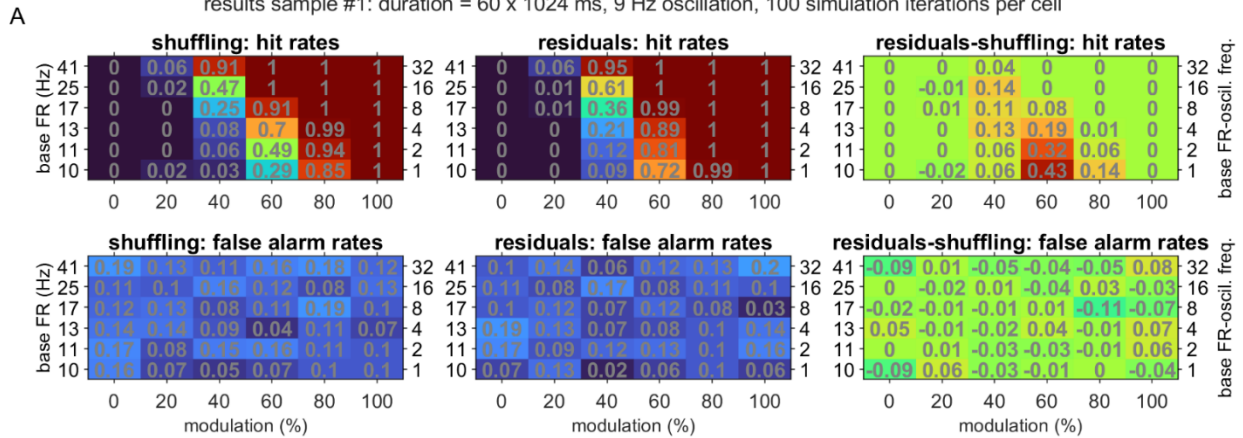


**Figure A12. Performance of the shuffling and residuals methods over a synthetic dataset generated using a short, absolute recovery period.** Panels (A)-(D): Plotting conventions, hit and false alarm definitions, and analysis procedures are identical to those described for the primary dataset depicted in Figure 4. Relative to the primary dataset, this secondary dataset differed in the use of an absolute, as opposed to relative RP ( $k = 0$ ), and a shorter RP duration ( $n_r = 3$  ms).



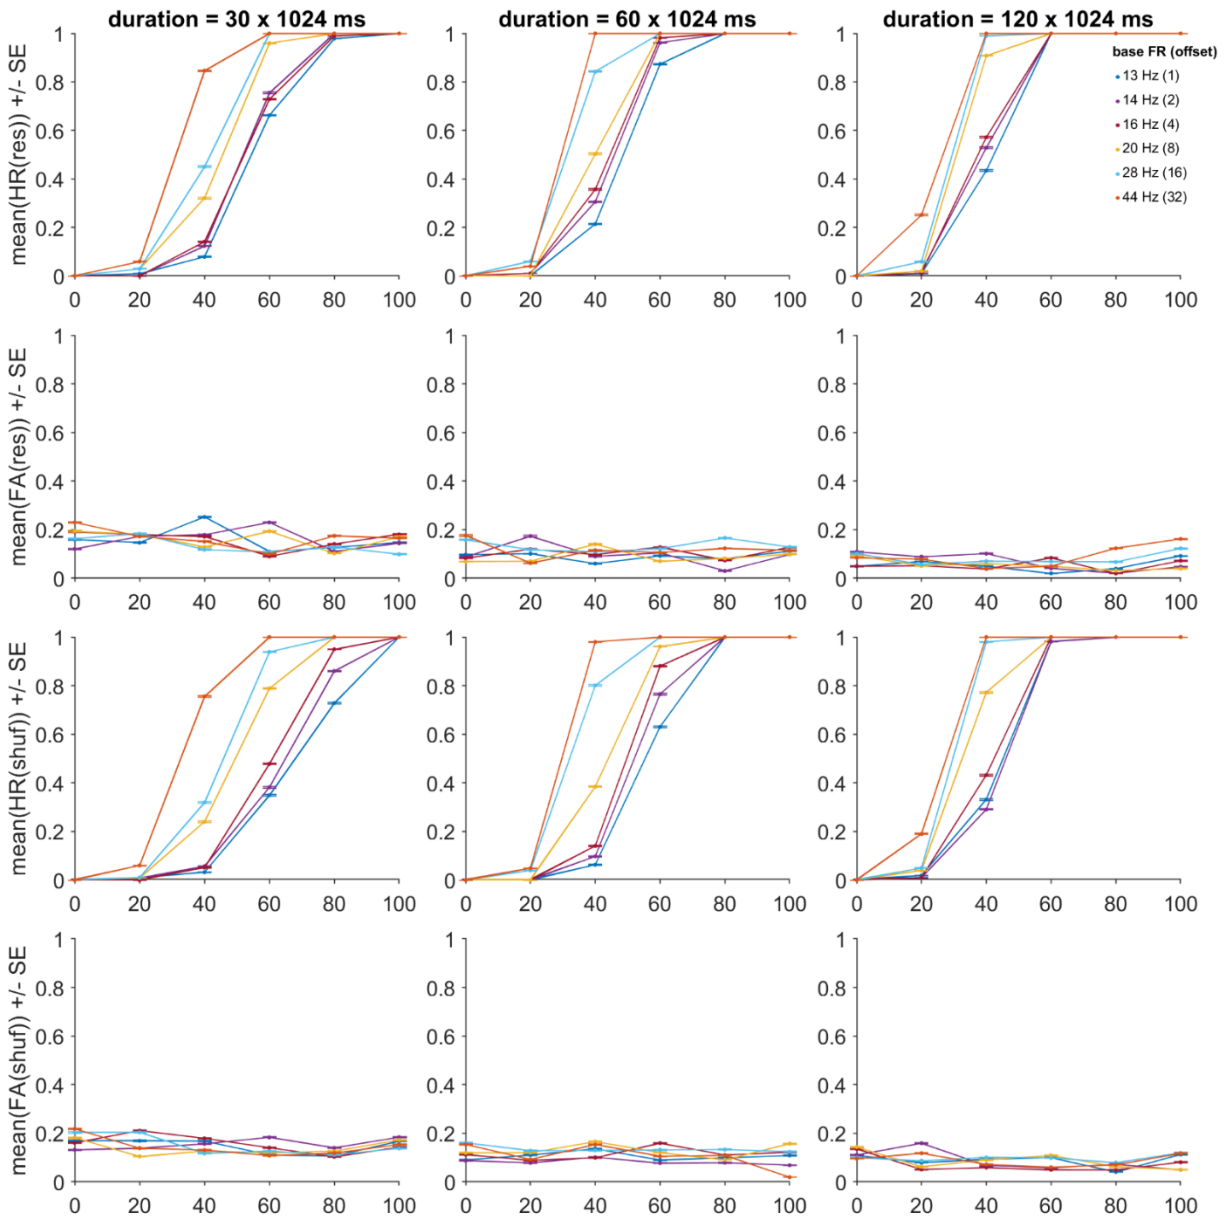
synthetic spike trains: method evaluation (low to moderate FRs, 18 ms RP,  $k = 0.7$ )

results sample #1: duration = 60 x 1024 ms, 9 Hz oscillation, 100 simulation iterations per cell

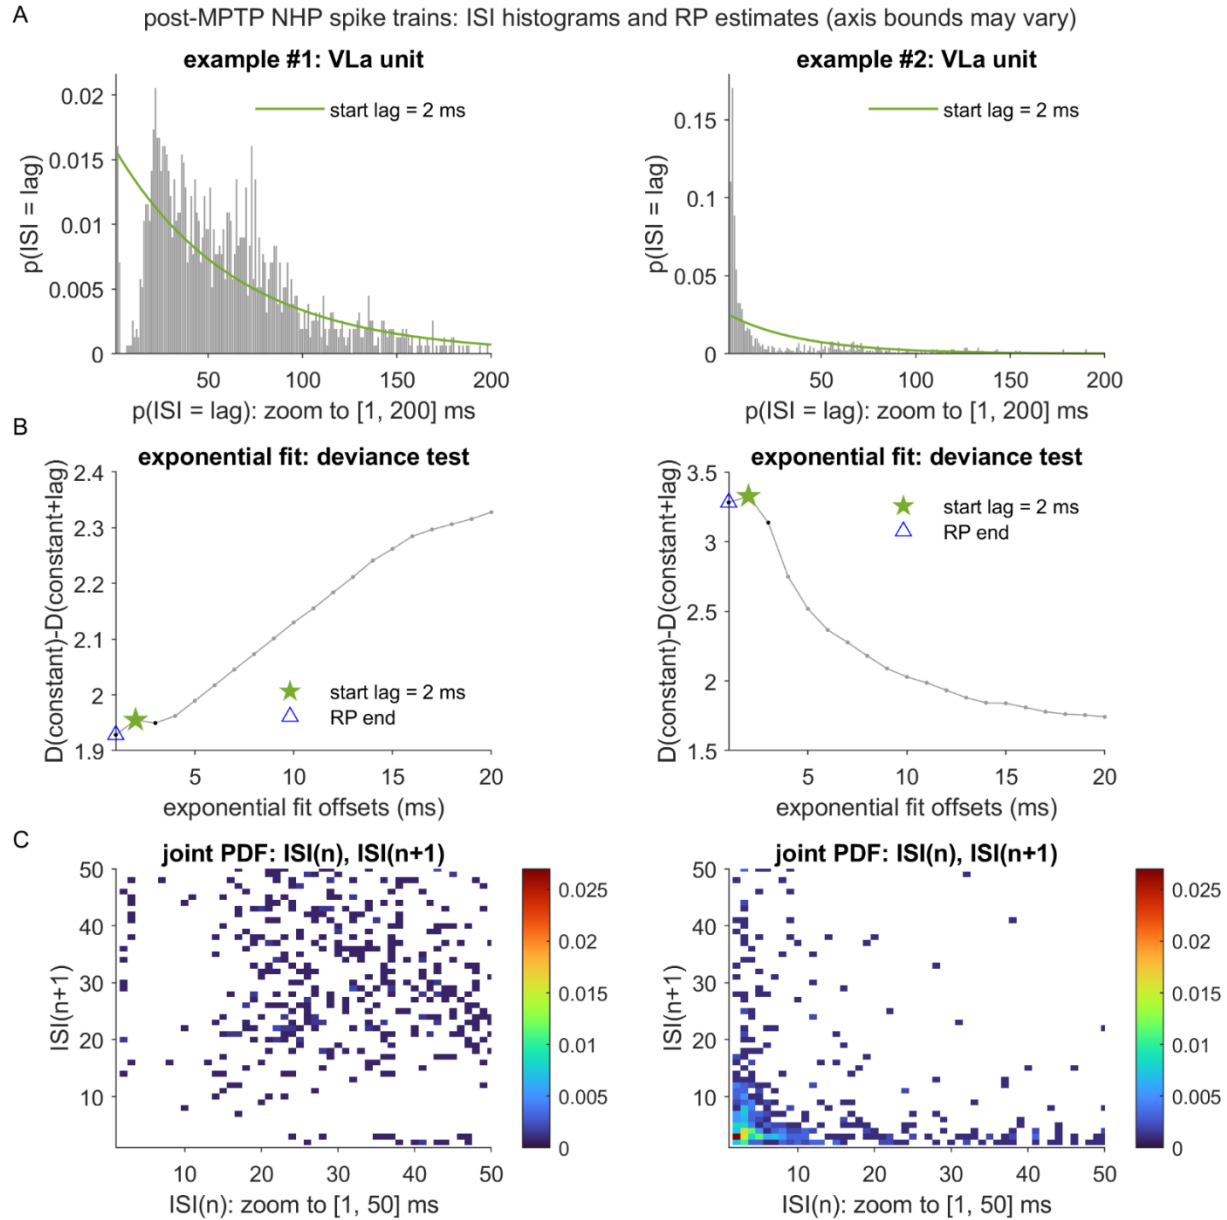


**Figure A14. Performance of the shuffling and residuals methods over a synthetic dataset generated using a lengthened relative recovery period.** Panels (A)-(D): Plotting conventions, hit and false alarm definitions, and analysis procedures are identical to those described for the primary dataset depicted in Figure 4. Relative to the primary dataset, this secondary dataset differed in the use of a longer relative RP ( $n_r = 18$  ms,  $k = 0.7$ ).

synthetic data: method hit rates and false alarms for  $f_{osc} = 12$  Hz (low to moderate FRs, 9 ms RP,  $k = 0.7$ )



**Figure A15. Residuals and shuffling hit and false alarm rates for the primary synthetic dataset, with the oscillation frequency ( $f_{osc}$ ) fixed at 12 Hz.** Abbreviations and plotting conventions follow from those described for Figure A8.-Figure A9..



**Figure A16. Inter-spike interval distributions and RP duration estimation for the two VLa units with suspected oscillatory and non-oscillatory features.** The left and right columns depict the same two empirical spike trains as were presented in Figure 11A-B. (A) Illustration of the procedure for obtaining the estimated RP duration ( $\hat{n}_r$ ) for the two spike trains. See Figure 3A-B and Chapter 3 for full details. For each unit, the figure highlights the exponential curve with the left anchor position (“start lag”) that returned the first local maximum in the curve goodness-of-fit plot. (B) Plots of the goodness-of-fit measure (the deviance difference statistic,  $\Delta D$ ) as a function of the first 20 starting positions of the exponential fits. Each  $\hat{n}_r$  estimate (blue triangles) was set equal to the post-spike lag immediately preceding the first local maximum in the corresponding  $\Delta D$  plot (green stars). (C) Joint probability density function (PDF) for all consecutive  $[ISI_n, ISI_{n+1}]$  inter-spike interval pairs. All remaining abbreviations follow from those described in Figure 3 and Figure 10.

## Appendix A.4 Supporting Tables

**Table A1. Parametric Results for the Methods' Hit Rate Differences: Low-to-Moderate Firing Rates.**

Table A1. Parametric Effects on Residuals-Shuffling Hit Rates				
Low-to-Moderate Firing Rates (9 ms RP, k = 0.7)				
Effect	Beta	SE	t	p
Intercept	0.084	0.00017	494.17	--
T	-0.00086	4.6e-06	-188.48	--
f <sub>osc</sub>	-0.0049	1.9e-05	-263.45	--
p <sub>base_offset</sub>	-0.0046	1.6e-05	-292.25	--
m	-0.17	0.0012	-143.68	--
m <sup>2</sup>	-0.48	0.0016	-296.41	--
T x f <sub>osc</sub>	3.4e-05	5e-07	68.25	--
T x p <sub>base_offset</sub>	3.5e-05	4.2e-07	82.42	--
f <sub>osc</sub> x p <sub>base_offset</sub>	0.00027	1.7e-06	155.21	--
T x m	-0.0024	3.1e-05	-77.28	--
f <sub>osc</sub> x m	-0.0028	0.00013	-21.75	7.6e-105
p <sub>base_offset</sub> x m	0.00057	0.00011	5.29	1.2e-07
T x m <sup>2</sup>	0.0049	4.4e-05	112.53	--
f <sub>osc</sub> x m <sup>2</sup>	0.03	0.00018	168.36	--
p <sub>base_offset</sub> x m <sup>2</sup>	0.029	0.00015	191.37	--
m x m <sup>2</sup>	1.1	0.0081	132.67	--
T x f <sub>osc</sub> x p <sub>base_offset</sub>	-9.9e-07	3e-08	-33.48	2e-245
T x f <sub>osc</sub> x m	8.4e-05	1.1e-06	73.92	--
T x p <sub>base_offset</sub> x m	6.4e-05	9.5e-07	66.88	--
f <sub>osc</sub> x p <sub>base_offset</sub> x m	0.00017	3.9e-06	44.01	--
T x f <sub>osc</sub> x m <sup>2</sup>	-0.00015	4.8e-06	-32.06	3.3e-225
T x p <sub>base_offset</sub> x m <sup>2</sup>	-0.00016	4e-06	-40.27	--
f <sub>osc</sub> x p <sub>base_offset</sub> x m <sup>2</sup>	-0.0015	1.7e-05	-90.54	--
T x m x m <sup>2</sup>	0.013	0.00022	60.29	--
f <sub>osc</sub> x m x m <sup>2</sup>	0.0075	0.00088	8.54	1.3e-17
p <sub>base_offset</sub> x m x m <sup>2</sup>	-0.012	0.00074	-16.60	7e-62
RP = recovery period; T = duration;				
f <sub>osc</sub> = oscillation frequency;				
p <sub>base_offset</sub> = offset of base firing rate from f <sub>osc</sub> ;				
m = oscillation modulation strength;				
-- = outside range of double representation				



**Table A2. Parametric Results for the Methods' False Alarm Rate Differences: Low-to-Moderate Firing Rates.**

Table A2. Parametric Effects on Residuals-Shuffling False Alarm				
Low-to-Moderate Firing Rates (9 ms RP, k = 0.7)				
Effect	Beta	SE	t	p
Intercept	-0.023	0.00017	-138.04	--
T	-0.00034	4.5e-06	-75.60	--
f <sub>osc</sub>	-0.0012	1.8e-05	-64.95	--
p <sub>base_offset</sub>	-4.3e-05	1.5e-05	-2.81	0.0049
m	-0.0026	0.0011	-2.28	0.023
m <sup>2</sup>	0.095	0.0016	59.43	--
T x f <sub>osc</sub>	1.2e-05	4.9e-07	24.87	1.9e-136
T x p <sub>base_offset</sub>	1.1e-05	4.1e-07	27.19	1.2e-162
f <sub>osc</sub> x p <sub>base_offset</sub>	4.3e-06	1.7e-06	2.54	0.011
T x m	0.001	3e-05	34.38	9.8e-259
f <sub>osc</sub> x m	-0.0031	0.00012	-25.25	1.4e-140
p <sub>base_offset</sub> x m	0.0055	0.0001	52.64	--
T x m <sup>2</sup>	0.00053	4.3e-05	12.32	6.9e-35
f <sub>osc</sub> x m <sup>2</sup>	-0.0027	0.00017	-15.26	1.4e-52
p <sub>base_offset</sub> x m <sup>2</sup>	0.0031	0.00015	21.28	2.1e-100
m x m <sup>2</sup>	0.031	0.0079	3.93	8.6e-05
T x f <sub>osc</sub> x p <sub>base_offset</sub>	-1e-06	2.9e-08	-34.66	7.9e-263
T x f <sub>osc</sub> x m	2.3e-05	1.1e-06	20.65	1e-94
T x p <sub>base_offset</sub> x m	2.1e-05	9.3e-07	22.26	9.5e-110
f <sub>osc</sub> x p <sub>base_offset</sub> x m	-0.00025	3.8e-06	-65.86	--
T x f <sub>osc</sub> x m <sup>2</sup>	2.1e-05	4.7e-06	4.39	1.1e-05
T x p <sub>base_offset</sub> x m <sup>2</sup>	-1.5e-05	3.9e-06	-3.72	0.0002
f <sub>osc</sub> x p <sub>base_offset</sub> x m <sup>2</sup>	-0.00043	1.6e-05	-26.54	4.8e-155
T x m x m <sup>2</sup>	-0.0034	0.00021	-16.24	2.8e-59
f <sub>osc</sub> x m x m <sup>2</sup>	0.0067	0.00086	7.73	1e-14
p <sub>base_offset</sub> x m x m <sup>2</sup>	-0.027	0.00073	-37.01	2.5e-299
RP = recovery period; T = duration;				
f <sub>osc</sub> = oscillation frequency;				
p <sub>base_offset</sub> = offset of base firing rate from f <sub>osc</sub> ;				
m = oscillation modulation strength;				
-- = outside range of double representation				

**Table A3. Parametric Results for the Methods' False Alarm Rate Differences: Low-to-Moderate Firing Rates.**

Table A3. Parametric Effects on Residuals-Shuffling Hit Rates Low-to-Moderate Firing Rates (9 ms RP, k = 0.7)				
Effect	Beta	SE	t	p
Intercept	0.1	0.0003	341.42	--
T	-0.00099	8e-06	-124.45	--
$\rho_{\text{base}}$	-0.0052	0.00012	-43.21	--
$\rho_{\text{base\_offset}}$	-0.008	0.00013	-59.81	--
m	-0.44	0.002	-217.82	--
$m^2$	-0.62	0.0029	-217.13	--
T x $\rho_{\text{base}}$	8.9e-05	3.2e-06	27.85	3.5e-170
T x $\rho_{\text{base\_offset}}$	4e-05	3.6e-06	11.29	1.5e-29
$\rho_{\text{base}}$ x $\rho_{\text{base\_offset}}$	0.00036	5.3e-05	6.66	2.7e-11
T x m	-0.004	5.4e-05	-72.62	--
$\rho_{\text{base}}$ x m	-0.018	0.00082	-22.29	7.4e-110
$\rho_{\text{base\_offset}}$ x m	0.013	0.00091	14.57	5e-48
T x $m^2$	0.0079	7.7e-05	103.32	--
$\rho_{\text{base}}$ x $m^2$	0.036	0.0011	31.79	6.4e-221
$\rho_{\text{base\_offset}}$ x $m^2$	0.053	0.0013	41.28	--
m x $m^2$	2.7	0.014	193.84	--
T x $\rho_{\text{base}}$ x $\rho_{\text{base\_offset}}$	-7.2e-06	9.2e-07	-7.81	5.8e-15
T x $\rho_{\text{base}}$ x m	2.2e-05	7.3e-06	3.03	0.0025
T x $\rho_{\text{base\_offset}}$ x m	0.00023	8.1e-06	28.19	3e-174
$\rho_{\text{base}}$ x $\rho_{\text{base\_offset}}$ x m	-4.1e-05	0.00012	-0.34	0.74
T x $\rho_{\text{base}}$ x $m^2$	-0.00045	3.1e-05	-14.64	1.8e-48
T x $\rho_{\text{base\_offset}}$ x $m^2$	-0.00062	3.4e-05	-17.96	5.3e-72
$\rho_{\text{base}}$ x $\rho_{\text{base\_offset}}$ x $m^2$	-0.0019	0.00051	-3.77	0.00016
T x m x $m^2$	0.024	0.00038	64.06	--
$\rho_{\text{base}}$ x m x $m^2$	0.11	0.0057	20.26	4e-91
$\rho_{\text{base\_offset}}$ x m x $m^2$	-0.084	0.0063	-13.21	7.7e-40
RP = recovery period; T = duration;				
$\rho_{\text{base}}$ = base firing rate;				
$\rho_{\text{base\_offset}}$ = offset of base firing rate from $f_{\text{osc}}$ ;				
m = oscillation modulation strength;				
-- = outside range of double representation				

**Table A4. Parametric Results for the False Alarm Differences: Manipulation of Base Firing Probability.**

Table A4. Parametric Effects on Residuals-Shuffling False Alarm				
Low-to-Moderate Firing Rates (9 ms RP, k = 0.7)				
Effect	Beta	SE	t	p
Intercept	-0.022	0.00033	-67.70	--
T	-0.00039	8.8e-06	-44.08	--
$p_{base}$	-0.0039	0.00013	-29.52	8.9e-191
$p_{base\_offset}$	-0.0036	0.00015	-24.56	6.4e-133
m	-0.018	0.0022	-8.17	3e-16
$m^2$	0.089	0.0032	28.06	1.2e-172
T x $p_{base}$	3.1e-05	3.5e-06	8.67	4.4e-18
T x $p_{base\_offset}$	1.7e-05	3.9e-06	4.23	2.3e-05
$p_{base}$ x $p_{base\_offset}$	0.0019	5.9e-05	32.70	1.8e-233
T x m	0.0019	6e-05	32.31	4.4e-228
$p_{base}$ x m	0.019	0.0009	21.40	1.9e-101
$p_{base\_offset}$ x m	-0.0023	0.001	-2.30	0.021
T x $m^2$	0.0011	8.4e-05	12.82	1.3e-37
$p_{base}$ x $m^2$	0.026	0.0013	20.63	2e-94
$p_{base\_offset}$ x $m^2$	0.022	0.0014	15.36	3.3e-53
m x $m^2$	0.25	0.016	16.26	2.2e-59
T x $p_{base}$ x $p_{base\_offset}$	-1.1e-05	1e-06	-10.86	1.9e-27
T x $p_{base}$ x m	6.7e-05	8e-06	8.43	3.5e-17
T x $p_{base\_offset}$ x m	0.00014	8.9e-06	15.81	3.2e-56
$p_{base}$ x $p_{base\_offset}$ x m	0.0015	0.00013	11.19	4.6e-29
T x $p_{base}$ x $m^2$	9.3e-05	3.4e-05	2.74	0.0061
T x $p_{base\_offset}$ x $m^2$	-0.00022	3.8e-05	-5.86	4.5e-09
$p_{base}$ x $p_{base\_offset}$ x $m^2$	-0.0058	0.00056	-10.21	1.8e-24
T x m x $m^2$	-0.0095	0.00042	-22.76	2.2e-114
$p_{base}$ x m x $m^2$	-0.11	0.0062	-17.11	1.4e-65
$p_{base\_offset}$ x m x $m^2$	0.054	0.007	7.78	7e-15
RP = recovery period; T = duration;				
$p_{base}$ = base firing rate;				
$p_{base\_offset}$ = offset of base firing rate from $f_{osc}$ ;				
m = oscillation modulation strength;				
-- = outside range of double representation				

**Table A5. Parametric Results for the Hit Rate Differences: High Firing Rates.**

Table A5. Parametric Effects on Residuals-Shuffling Hit Rates High Firing Rates (9 ms RP, k = 0.7)				
Effect	Beta	SE	t	p
Intercept	0.0011	6.2e-05	18.48	3.4e-76
T	-4.8e-05	1.7e-06	-28.85	6.6e-183
f <sub>osc</sub>	-0.00011	6.8e-06	-15.83	1.9e-56
p <sub>base_offset</sub>	-0.00019	3.6e-06	-51.56	--
m	-0.018	0.00042	-42.45	--
m <sup>2</sup>	0.045	0.00059	75.82	--
T x f <sub>osc</sub>	2.6e-06	1.8e-07	14.60	2.8e-48
T x p <sub>base_offset</sub>	5.1e-06	9.7e-08	52.69	--
f <sub>osc</sub> x p <sub>base_offset</sub>	5.1e-06	4e-07	12.96	2.2e-38
T x m	0.00056	1.1e-05	49.35	--
f <sub>osc</sub> x m	0.00094	4.6e-05	20.44	8.3e-93
p <sub>base_offset</sub> x m	0.0017	2.5e-05	69.09	--
T x m <sup>2</sup>	-0.00045	1.6e-05	-28.05	5.7e-173
f <sub>osc</sub> x m <sup>2</sup>	0.0013	6.5e-05	20.03	3e-89
p <sub>base_offset</sub> x m <sup>2</sup>	0.0018	3.5e-05	51.76	--
m x m <sup>2</sup>	-0.015	0.0029	-5.27	1.4e-07
T x f <sub>osc</sub> x p <sub>base_offset</sub>	-1.5e-08	6.8e-09	-2.22	0.026
T x f <sub>osc</sub> x m	5.5e-08	4.1e-07	0.13	0.89
T x p <sub>base_offset</sub> x m	4.5e-06	2.2e-07	20.60	3.3e-94
f <sub>osc</sub> x p <sub>base_offset</sub> x m	-7.2e-06	9e-07	-8.00	1.3e-15
T x f <sub>osc</sub> x m <sup>2</sup>	-2.5e-05	1.7e-06	-14.19	1.1e-45
T x p <sub>base_offset</sub> x m <sup>2</sup>	-6e-05	9.3e-07	-64.42	--
f <sub>osc</sub> x p <sub>base_offset</sub> x m <sup>2</sup>	-2.8e-05	3.8e-06	-7.26	4e-13
T x m x m <sup>2</sup>	-0.0017	7.8e-05	-22.31	3.2e-110
f <sub>osc</sub> x m x m <sup>2</sup>	-0.0078	0.00032	-24.22	1.6e-129
p <sub>base_offset</sub> x m x m <sup>2</sup>	-0.013	0.00017	-74.32	--
RP = recovery period; T = duration;				
f <sub>osc</sub> = oscillation frequency;				
p <sub>base_offset</sub> = offset of base firing rate from f <sub>osc</sub> ;				
m = oscillation modulation strength;				
-- = outside range of double representation				

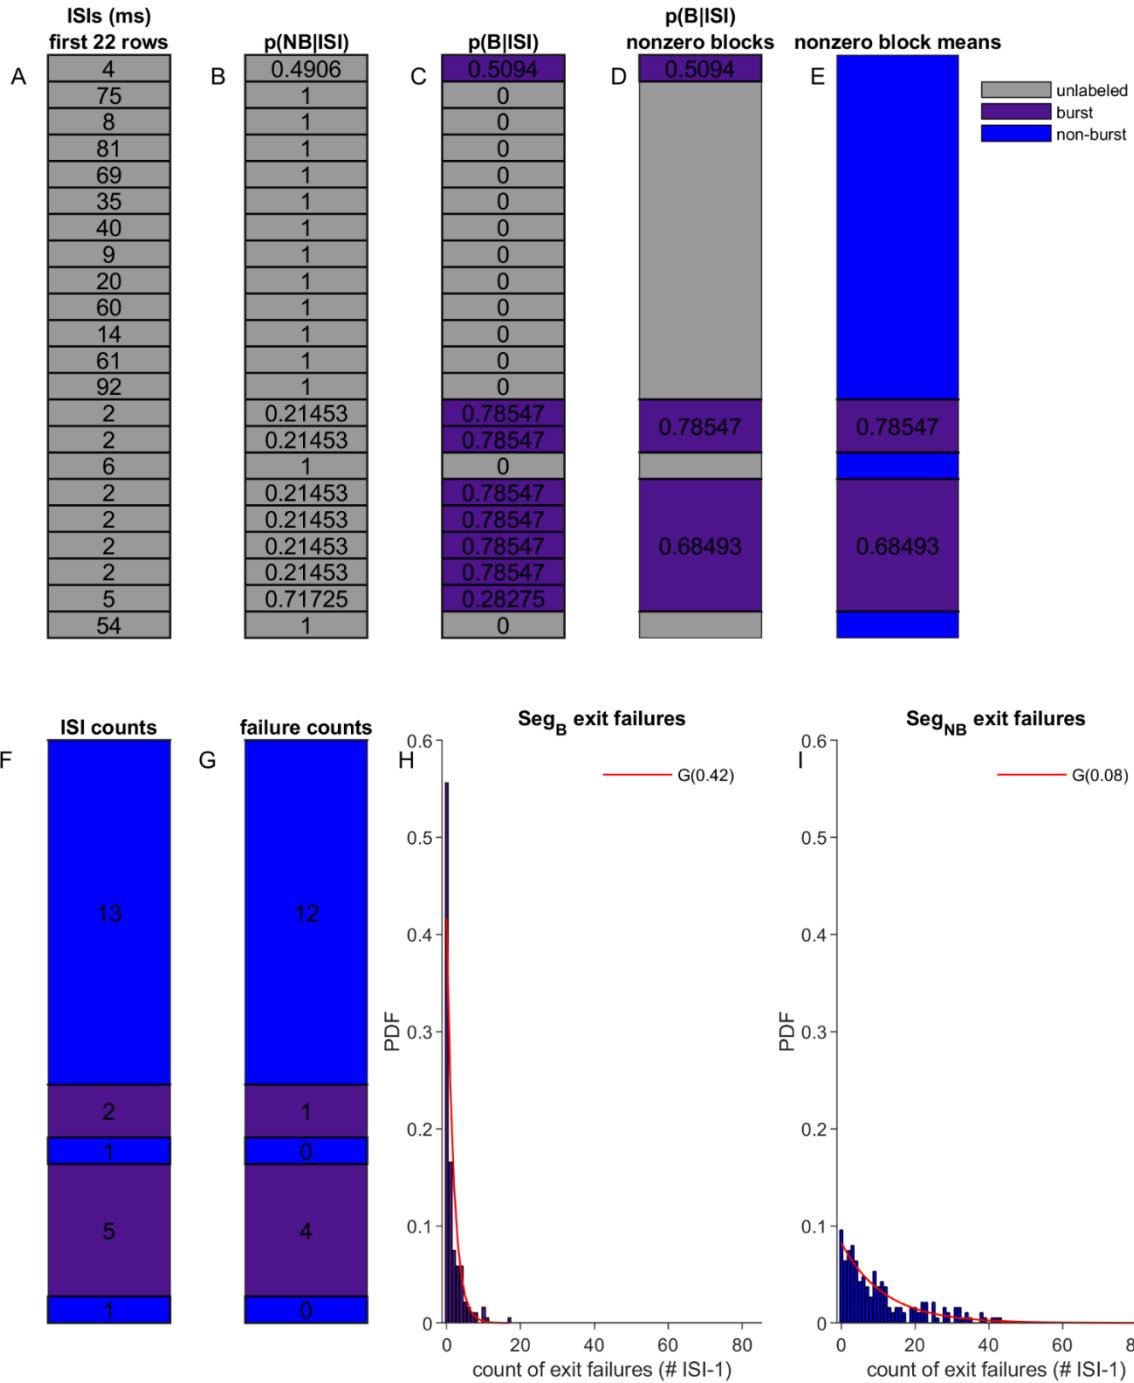
**Table A6. Parametric Results for the False Alarm Differences: High Firing Rates.**

Table A6. Parametric Effects on Residuals-Shuffling False Alarm High Firing Rates (9 ms RP, k = 0.7)				
Effect	Beta	SE	t	p
Intercept	-0.014	0.00022	-63.65	--
T	1.5e-05	6e-06	2.53	0.011
f <sub>osc</sub>	-0.00052	2.4e-05	-21.37	3e-101
p <sub>base_offset</sub>	0.00024	1.3e-05	18.24	2.5e-74
m	0.12	0.0015	76.88	--
m <sup>2</sup>	0.23	0.0021	107.57	--
T x f <sub>osc</sub>	-5.4e-06	6.5e-07	-8.25	1.6e-16
T x p <sub>base_offset</sub>	-6.7e-06	3.5e-07	-19.16	8.6e-82
f <sub>osc</sub> x p <sub>base_offset</sub>	1.3e-05	1.4e-06	9.33	1e-20
T x m	4.2e-05	4.1e-05	1.04	0.3
f <sub>osc</sub> x m	-0.0089	0.00017	-53.48	--
p <sub>base_offset</sub> x m	0.002	8.9e-05	22.75	1.6e-114
T x m <sup>2</sup>	0.0005	5.7e-05	8.83	1e-18
f <sub>osc</sub> x m <sup>2</sup>	-0.079	0.00023	-336.62	--
p <sub>base_offset</sub> x m <sup>2</sup>	0.0057	0.00013	45.72	--
m x m <sup>2</sup>	-0.11	0.011	-10.32	5.9e-25
T x f <sub>osc</sub> x p <sub>base_offset</sub>	-6.5e-07	2.4e-08	-26.43	7.2e-154
T x f <sub>osc</sub> x m	-0.00026	1.5e-06	-178.70	--
T x p <sub>base_offset</sub> x m	2.8e-05	7.9e-07	35.75	1.9e-279
f <sub>osc</sub> x p <sub>base_offset</sub> x m	-0.00059	3.2e-06	-181.70	--
T x f <sub>osc</sub> x m <sup>2</sup>	-0.00069	6.2e-06	-110.94	--
T x p <sub>base_offset</sub> x m <sup>2</sup>	0.00015	3.3e-06	45.24	--
f <sub>osc</sub> x p <sub>base_offset</sub> x m <sup>2</sup>	-0.0016	1.4e-05	-116.58	--
T x m x m <sup>2</sup>	0.006	0.00028	21.19	1.3e-99
f <sub>osc</sub> x m x m <sup>2</sup>	-0.15	0.0012	-129.27	--
p <sub>base_offset</sub> x m x m <sup>2</sup>	0.0062	0.00062	10.05	9.2e-24
RP = recovery period; T = duration;				
f <sub>osc</sub> = oscillation frequency;				
p <sub>base_offset</sub> = offset of base firing rate from f <sub>osc</sub> ;				
m = oscillation modulation strength;				
-- = outside range of double representation				

## **Appendix B Supporting Figures and Tables for the Two-State Residuals Method**

### **Appendix B.1 Supporting Figures**

Schematic: Initialization of State Transition Matrix

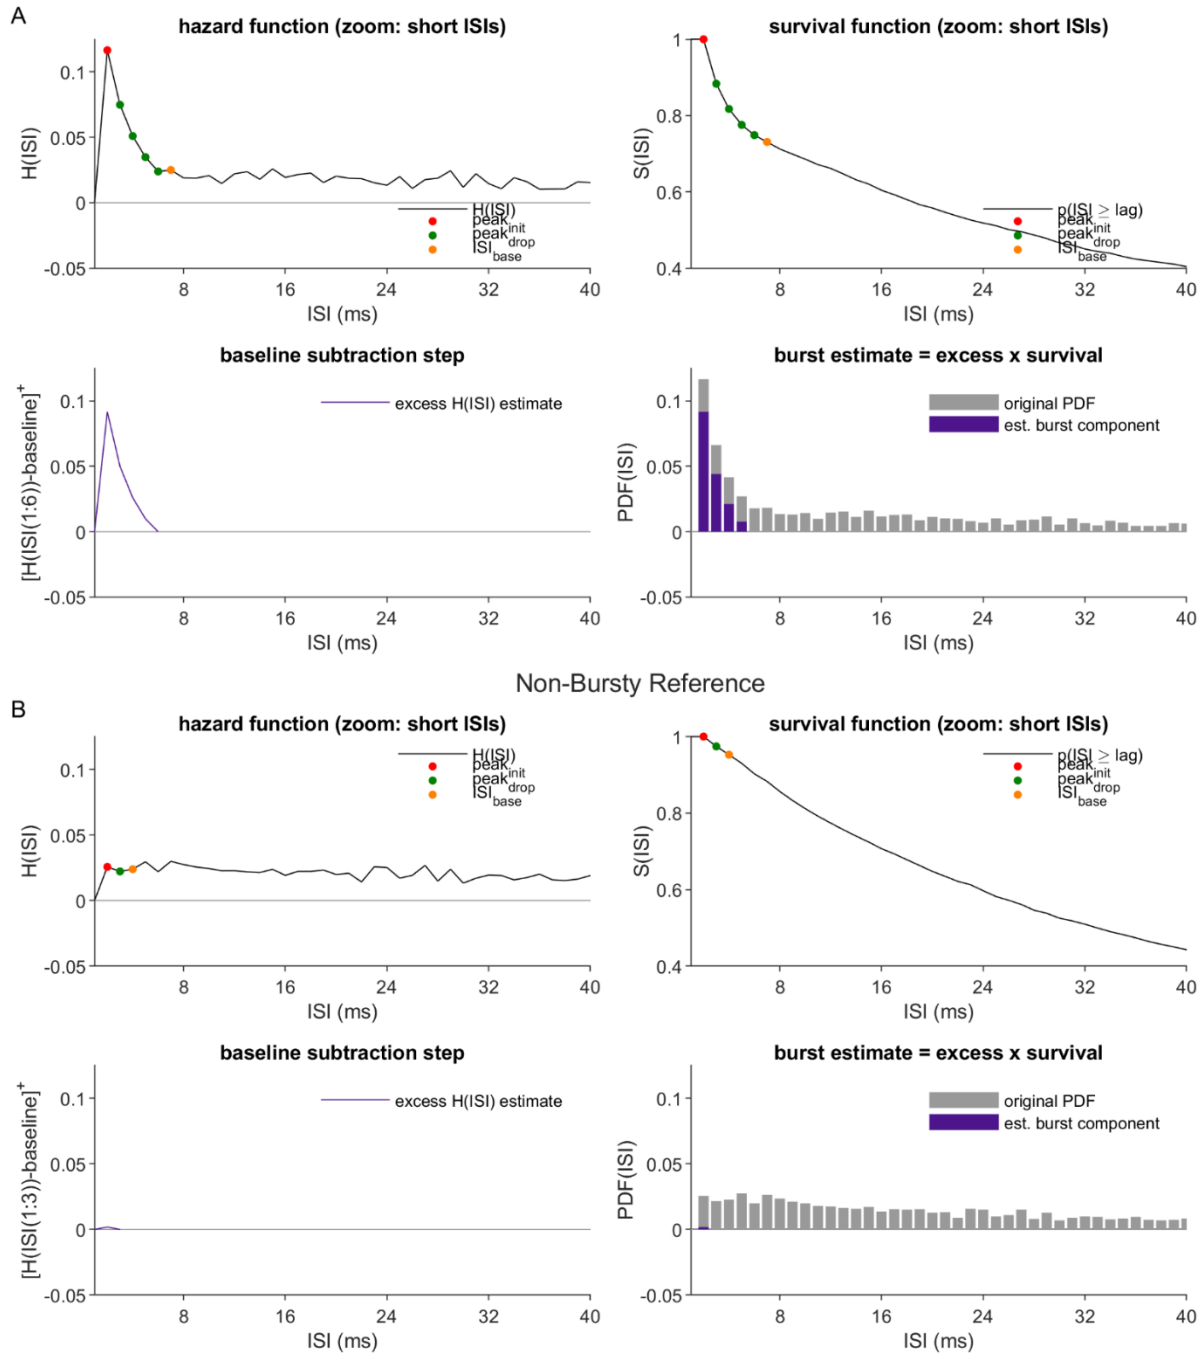


**Figure B1.** Illustration of the procedure used to obtain an initial estimate of the state transition matrix, **A**, to be applied to the initialization of the Hidden Markov Model (HMM). See Section 5.3.3 for a detailed description. All plots were derived from the same example synthetic unit as was introduced in Section 6.2. (A) The algorithm begins by generating the list of inter-spike intervals for the spike train. The illustrations in Panels A-G are limited to the first 22 ISIs (actual  $n = 2736$ ). The second step entails the estimation, for each  $ISI_i$  in the list, of the probabilities that  $ISI_i$  arose from either the non-burst (NB) state (panel B) or

burst state (panel C), given the prior probability distribution over ISIs (i.e., the observed ISI PDF), the initial estimates of the conditional distributions  $P(\text{ISI}|\text{NB})$  and  $P(\text{ISI}|\text{B})$ , and the prior probability of the non-burst and burst states' occurrence ( $p_{\text{NB}}$ ,  $p_{\text{B}}$ ), as obtained through the procedures described in Sections 5.3.1-5.3.2. The result of the second step is a pair of vectors, representing the posterior likelihood estimates for  $P(\text{NB}|\text{ISI}_i)$  and  $P(\text{B}|\text{ISI}_i)$ , as obtained through Bayes' rule. Using the vector of  $P(\text{B}|\text{ISI})$  values, the algorithm identifies a preliminary set of candidate "B" segments (purple shading) by selecting all the contiguous runs of nonzero entries. This liberal initial estimate was filtered down to a more conservative one through a two step process. (D) First, each candidate "B" segment was assigned a weight equal to the mean of the  $P(\text{B}|\text{ISI})$  values contained in that segment. (E) Only the most highly weighted segments retained their "burst" labels. The cutoff was set to admit as many of the top-weighted segments as possible, while ensuring that the total count of ISIs represented across those segments did not exceed the expected count of burst-associated ISIs, given  $\hat{p}_{\text{B}}$ . All remaining segments (including the rejected candidate "B" segments) were assigned an "NB" label. (F) Representation of the ISI counts for each segment, and (G) the same values minus one, corresponding to the number of spikes that failed to trigger an exit to the next state. Note that the proposed HMM assumes that the per-state failure counts will be geometrically distributed. (G) and (H) depict the fits of geometric distributions to the failure count data. The estimated  $p$  parameters for these distributions are used to set the off-diagonal elements of the initial transition matrix  $A$ .



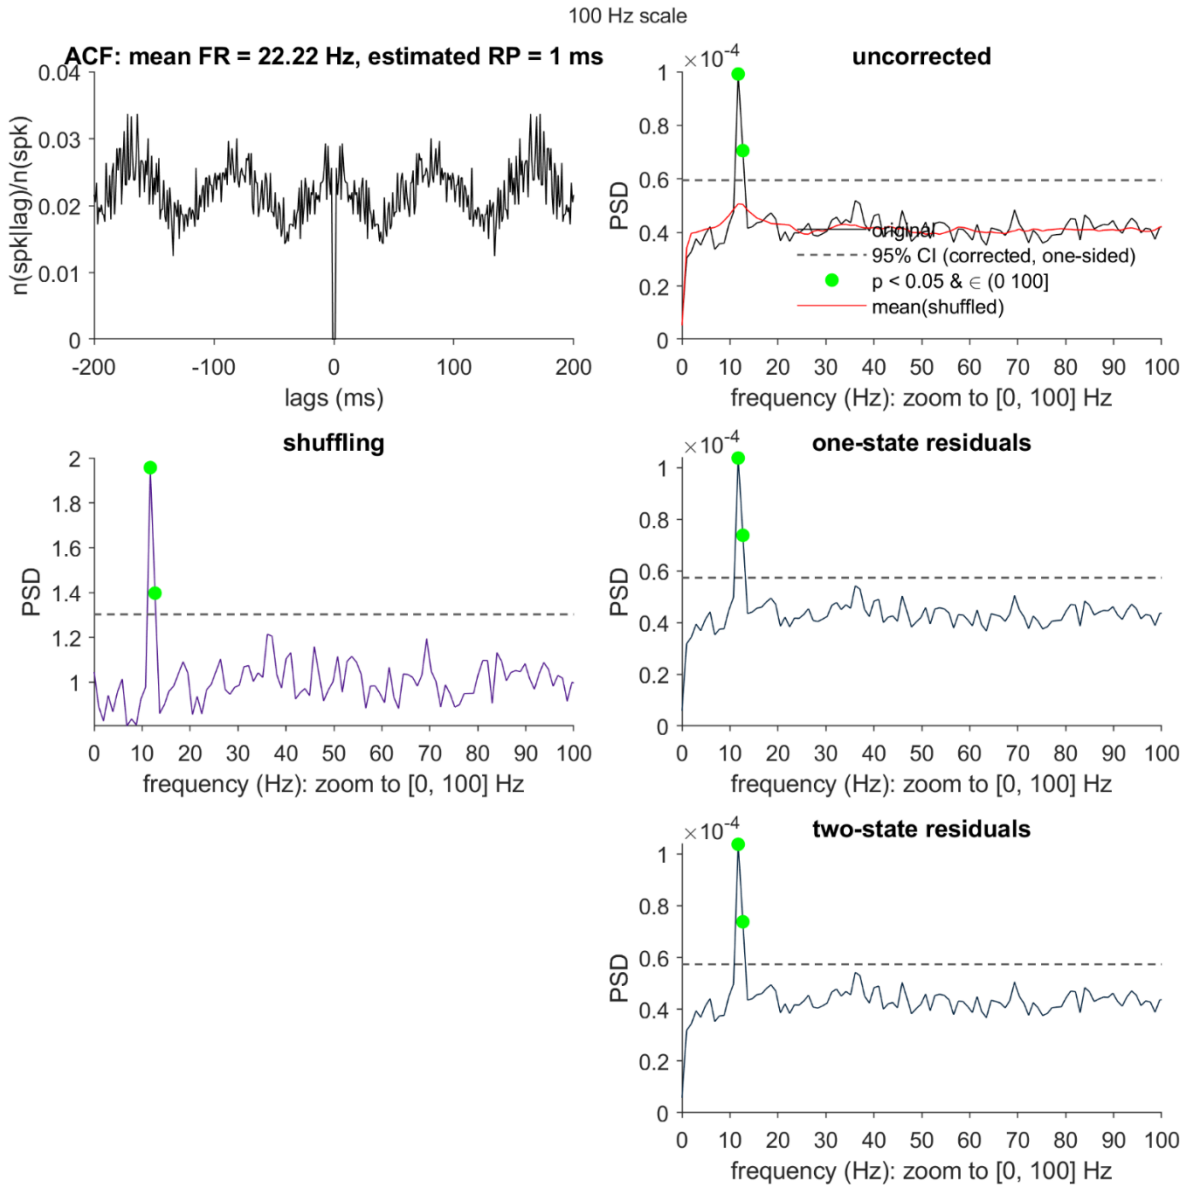
Synthetic Approximation to Fig. A11B Unit: Decomposition into Two ISI Distributions  
Bursty Spike Train



**Figure B2. Demonstration of the ISI separation procedure, for the simulations informed by the empirical unit in Figure 11B.** (A) Application to the “bursty” simulated unit. Between the burst- and non-burst states, only the base FR differed (RP = 1 ms absolute). As expected, the splitting procedure divides the early-ISI zone of the original histogram across the non-burst and burst states. (B) Application to the “non-bursty” reference unit. Although the estimated “non-burst” distribution is mostly but not

entirely empty (see  $ISI = 2$  ms). This nonzero mass is sufficient for the unit to be categorized as “bursty,” but the negligible mass assigned to the burst distribution is factored into downstream stages in the burst detection pipeline.

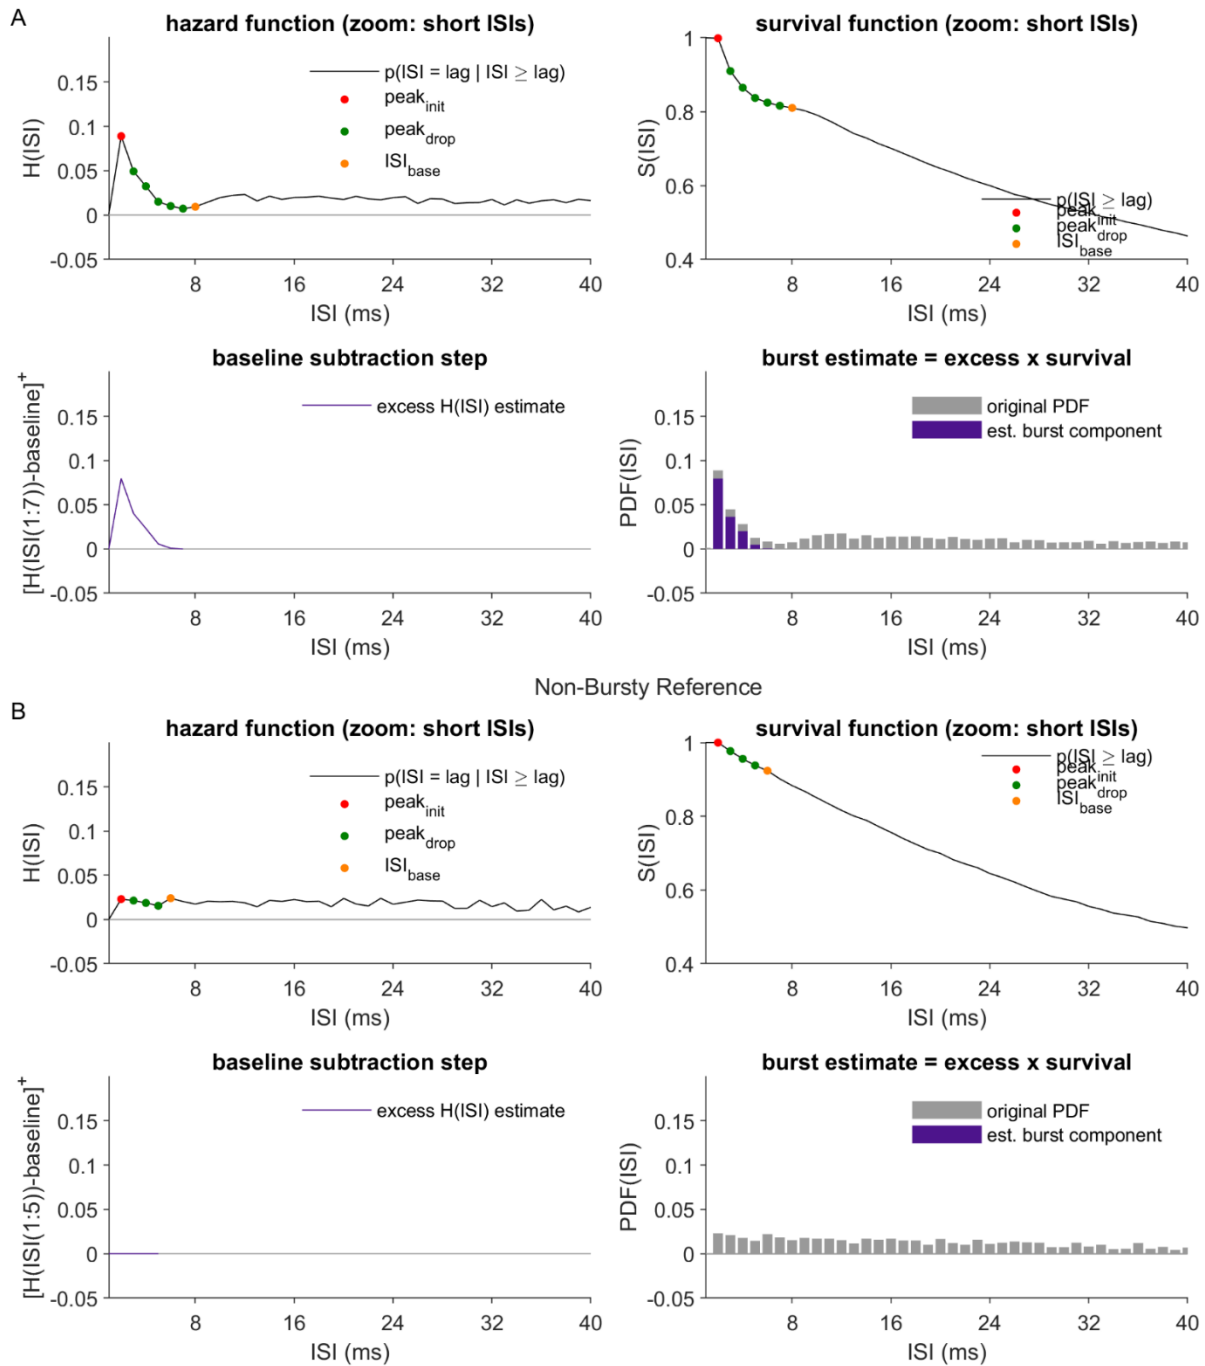
Synthetic Approximation to Fig. A11B Unit (Non-Bursty): Comparison of Corrected PSDs



**Figure B3. Comparison of the shuffling, one-state (1S) residuals output and two state (2S) residuals output for the non-bursty control unit, matched to the simulation designed to approximate the empirical unit from Figure 11B.**

Formatting conventions and abbreviations follow from those used in Figure 10.

Synthetic Unit with a Moderate-Length Non-burst RP: Decomposition into Two ISI Distributions  
Bursty Spike Train



**Figure B4.** Demonstration of the ISI separation procedure, for parameter settings expected to generate a rise in the non-burst hazard function over the shortest ISIs (for the “bursty” unit only). (A) Application to the “bursty” simulated unit. (B) Application to the “non-bursty” reference unit. Abbreviations and formatting conventions follow those from Figure B2.

## Bibliography

- [1] E. R. Kandel, J. D. Koester, S. H. Mack, and S. A. Siegelbaum, *Principles of Neural Science, Sixth Edition*. McGraw-Hill Education, 2021.
- [2] A. Matzner and I. Bar-Gad, "Quantifying spike train oscillations: biases, distortions and solutions," *PLoS Comput Biol*, vol. 11, no. 4, p. e1004252, Apr 2015, doi: 10.1371/journal.pcbi.1004252.
- [3] J. E. Rubin, "Computational models of basal ganglia dysfunction: the dynamics is in the details," *Curr Opin Neurobiol*, vol. 46, pp. 127-135, Oct 2017, doi: 10.1016/j.conb.2017.08.011.
- [4] G. Buzsaki, C. A. Anastassiou, and C. Koch, "The origin of extracellular fields and currents--EEG, ECoG, LFP and spikes," *Nat Rev Neurosci*, vol. 13, no. 6, pp. 407-20, May 18 2012, doi: 10.1038/nrn3241.
- [5] S. Little and P. Brown, "Debugging Adaptive Deep Brain Stimulation for Parkinson's Disease," *Mov Disord*, vol. 35, no. 4, pp. 555-561, Apr 2020, doi: 10.1002/mds.27996.
- [6] R. E. Kass, U. T. Eden, and E. N. Brown, *Analysis of Neural Data*. Springer New York, 2014.
- [7] M. Rivlin-Etzion, Y. Ritov, G. Heimer, H. Bergman, and I. Bar-Gad, "Local shuffling of spike trains boosts the accuracy of spike train spectral analysis," *J Neurophysiol*, vol. 95, no. 5, pp. 3245-56, May 2006, doi: 10.1152/jn.00055.2005.
- [8] W. Truccolo, U. T. Eden, M. R. Fellows, J. P. Donoghue, and E. N. Brown, "A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects," *Journal of neurophysiology*, vol. 93, no. 2, pp. 1074-1089, 2005.
- [9] W. Bair, C. Koch, W. Newsome, and K. Britten, "Power spectrum analysis of bursting cells in area MT in the behaving monkey," *J Neurosci*, vol. 14, no. 5 Pt 1, pp. 2870-92, May 1994. [Online]. Available: <https://www.ncbi.nlm.nih.gov/pubmed/8182445>.
- [10] B. W. Edwards, G. H. Wakefield, and N. L. Powers, "The spectral shaping of neural discharges by refractory effects," *J Acoust Soc Am*, vol. 93, no. 6, pp. 3353-64, Jun 1993, doi: 10.1121/1.405718.
- [11] K. Arai and R. E. Kass, "Inferring oscillatory modulation in neural spike trains," *PLoS Comput Biol*, vol. 13, no. 10, p. e1005596, Oct 2017, doi: 10.1371/journal.pcbi.1005596.
- [12] S. V. Sarma, M. L. Cheng, U. Eden, Z. Williams, E. N. Brown, and E. Eskandar, "The effects of cues on neurons in the basal ganglia in Parkinson's disease," *Front Integr Neurosci*, vol. 6, p. 40, 2012, doi: 10.3389/fnint.2012.00040.
- [13] S. Saxena, S. V. Sarma, S. R. Patel, S. Santaniello, E. N. Eskandar, and J. T. Gale, "Modulations in Oscillatory Activity of Globus Pallidus Internus Neurons During a Directed Hand Movement Task-A Primary Mechanism for Motor Planning," *Front Syst Neurosci*, vol. 13, p. 15, 2019, doi: 10.3389/fnsys.2019.00015.
- [14] E. Cotterill, P. Charlesworth, C. W. Thomas, O. Paulsen, and S. J. Eglen, "A comparison of computational methods for detecting bursts in neuronal spike trains and their application

- to human stem cell-derived neuronal networks," *J Neurophysiol*, vol. 116, no. 2, pp. 306-21, Aug 1 2016, doi: 10.1152/jn.00093.2016.
- [15] Z. Friedenberger, E. Harkin, K. Tóth, and R. Naud, "Silences, spikes and bursts: Three-part knot of the neural code," *The Journal of Physiology*, 2023.
- [16] B. Ermentrout and D. H. Terman, *Mathematical foundations of neuroscience*. Springer, 2010.
- [17] F. Grüneis, M. Nakao, and M. Yamamoto, "Counting statistics of 1/f fluctuations in neuronal spike trains," *Biological cybernetics*, vol. 62, no. 5, pp. 407-413, 1990.
- [18] S. Tokdar, P. Xi, R. C. Kelly, and R. E. Kass, "Detection of bursts in extracellular spike trains using hidden semi-Markov point process models," *J Comput Neurosci*, vol. 29, no. 1-2, pp. 203-212, Aug 2010, doi: 10.1007/s10827-009-0182-2.
- [19] K. M. Cox, D. Kase, T. Znati, and R. S. Turner, "Detecting rhythmic spiking through the power spectra of point process model residuals," *bioRxiv*, p. 2023.09.08.556120, 2023.
- [20] A. P. Buccino *et al.*, "SpikeInterface, a unified framework for spike sorting," *Elife*, vol. 9, Nov 10 2020, doi: 10.7554/eLife.61834.
- [21] G. Buzsáki, *Rhythms of the Brain*. Oxford university press, 2006.
- [22] L. Squire, D. Berg, F. E. Bloom, S. Du Lac, A. Ghosh, and N. C. Spitzer, *Fundamental neuroscience*. Academic press, 2012.
- [23] D. J. Surmeier, J. N. Mercer, and C. S. Chan, "Autonomous pacemakers in the basal ganglia: who needs excitatory synapses anyway?," *Curr Opin Neurobiol*, vol. 15, no. 3, pp. 312-8, Jun 2005, doi: 10.1016/j.conb.2005.05.007.
- [24] M. Bachetti and D. Huppenkothen, "No time for dead time: use the Fourier amplitude differences to normalize dead-time-affected periodograms," *The Astrophysical Journal Letters*, vol. 853, no. 2, p. L21, 2018.
- [25] W. Zhang, K. Jahoda, J. Swank, E. Morgan, and A. Giles, "Dead-time modifications to fast Fourier transform power spectra," *Astrophysical Journal v. 449, p. 930*, vol. 449, p. 930, 1995.
- [26] D. M. Halliday, J. R. Rosenberg, A. M. Amjad, P. Breeze, B. A. Conway, and S. F. Farmer, "A framework for the analysis of mixed time series/point process data--theory and application to the study of physiological tremor, single motor unit discharges and electromyograms," *Prog Biophys Mol Biol*, vol. 64, no. 2-3, pp. 237-78, 1995, doi: 10.1016/s0079-6107(96)00009-0.
- [27] R. Durrett, *Essentials of Stochastic Processes*. Springer International Publishing, 2016.
- [28] W. Gerstner, W. M. Kistler, R. Naud, and L. Paninski, *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press, 2014.
- [29] T. C. Whalen, A. M. Willard, J. E. Rubin, and A. H. Gittis, "Delta oscillations are a robust biomarker of dopamine depletion severity and motor dysfunction in awake mice," *J Neurophysiol*, vol. 124, no. 2, pp. 312-329, Aug 1 2020, doi: 10.1152/jn.00158.2020.
- [30] D. J. Thomson, "Spectrum estimation and harmonic analysis," *Proceedings of the IEEE*, vol. 70, no. 9, pp. 1055-1096, 1982.
- [31] M. K. Baaske *et al.*, "Parkinson's disease uncovers an underlying sensitivity of subthalamic nucleus neurons to beta-frequency cortical input in vivo," *Neurobiol Dis*, vol. 146, p. 105119, Dec 2020, doi: 10.1016/j.nbd.2020.105119.

- [32] T. Hasegawa, S. Chiken, K. Kobayashi, and A. Nambu, "Subthalamic nucleus stabilizes movements by reducing neural spike variability in monkey basal ganglia," *Nat Commun*, vol. 13, no. 1, p. 2233, Apr 25 2022, doi: 10.1038/s41467-022-29750-2.
- [33] T. Hashimoto *et al.*, "Neuronal activity and outcomes from thalamic surgery for spinocerebellar ataxia," *Ann Clin Transl Neurol*, vol. 5, no. 1, pp. 52-63, Jan 2018, doi: 10.1002/acn3.508.
- [34] B. Pasquereau and R. S. Turner, "Primary motor cortex of the parkinsonian monkey: differential effects on the spontaneous activity of pyramidal tract-type neurons," *Cereb Cortex*, vol. 21, no. 6, pp. 1362-78, Jun 2011, doi: 10.1093/cercor/bhq217.
- [35] Y. Tachibana, H. Iwamuro, H. Kita, M. Takada, and A. Nambu, "Subthalamo-pallidal interactions underlying parkinsonian neuronal oscillations in the primate basal ganglia," *Eur J Neurosci*, vol. 34, no. 9, pp. 1470-84, Nov 2011, doi: 10.1111/j.1460-9568.2011.07865.x.
- [36] S. Sani, S. Shimamoto, R. S. Turner, N. Levesque, and P. A. Starr, "Microelectrode recording in the posterior hypothalamic region in humans," *Neurosurgery*, vol. 64, no. 3 Suppl, pp. ons161-7; discussion ons167-9, Mar 2009, doi: 10.1227/01.NEU.0000334051.91501.E3.
- [37] S. L. Schmidt, E. Y. Chew, D. V. Bennett, M. A. Hammad, and F. Frohlich, "Differential effects of cholinergic and noradrenergic neuromodulation on spontaneous cortical network dynamics," *Neuropharmacology*, vol. 72, pp. 259-73, Sep 2013, doi: 10.1016/j.neuropharm.2013.04.045.
- [38] W. Van Drongelen, *Signal processing for neuroscientists*. Academic press, 2018.
- [39] C. R. Gallistel, E. Rolls, and D. Greene, "Neuron function inferred from behavioral and electrophysiological estimates of refractory period," *Science*, vol. 166, no. 3908, pp. 1028-30, Nov 21 1969, doi: 10.1126/science.166.3908.1028.
- [40] P. Kara, P. Reinagel, and R. C. Reid, "Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons," *Neuron*, vol. 27, no. 3, pp. 635-46, Sep 2000, doi: 10.1016/s0896-6273(00)00072-6.
- [41] R. E. Kass and V. Ventura, "A spike-train probability model," *Neural Comput*, vol. 13, no. 8, pp. 1713-20, Aug 2001, doi: 10.1162/08997660152469314.
- [42] E. M. Izhikevich, "Neural excitability, spiking and bursting," *International journal of bifurcation and chaos*, vol. 10, no. 06, pp. 1171-1266, 2000.
- [43] N. Technologies. "NeuroExplorer Version 5 User Manual." <https://www.neuroexplorer.com/downloads/NeuroExplorer5Manual.pdf> (accessed).
- [44] V. Pasquale, S. Martinoia, and M. Chiappalone, "A self-adapting approach for the detection of bursts and network bursts in neuronal cultures," *Journal of computational neuroscience*, vol. 29, pp. 213-229, 2010.
- [45] C. R. Legendy and M. Salcman, "Bursts and recurrences of bursts in the spike trains of spontaneously active striate cortex neurons," *J Neurophysiol*, vol. 53, no. 4, pp. 926-39, Apr 1985, doi: 10.1152/jn.1985.53.4.926.
- [46] B. Gourévitch and J. J. Eggermont, "A nonparametric approach for detection of bursts in spike trains," *Journal of neuroscience methods*, vol. 160, no. 2, pp. 349-358, 2007.
- [47] D. Ko, C. J. Wilson, C. J. Lobb, and C. A. Paladini, "Detection of bursts and pauses in spike trains," *J Neurosci Methods*, vol. 211, no. 1, pp. 145-58, Oct 15 2012, doi: 10.1016/j.jneumeth.2012.08.013.

- [48] C. M. Bishop and N. M. Nasrabadi, *Pattern recognition and machine learning* (no. 4). Springer, 2006.
- [49] K. P. Murphy, *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [50] L. R. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proceedings of the IEEE*, vol. 77, no. 2, pp. 257-286, 1989.
- [51] N. A. Heckert and J. J. Filliben, "Nist/sematech e-handbook of statistical methods; chapter 1: Exploratory data analysis," 2003.
- [52] J. Pohle, R. Langrock, F. M. Van Beest, and N. M. Schmidt, "Selecting the number of states in hidden Markov models: pragmatic solutions illustrated using animal movement," *Journal of Agricultural, Biological and Environmental Statistics*, vol. 22, pp. 270-293, 2017.
- [53] I. Bar-Gad, Y. Ritov, and H. Bergman, "The neuronal refractory period causes a short-term peak in the autocorrelation function," *J Neurosci Methods*, vol. 104, no. 2, pp. 155-63, Jan 15 2001, doi: 10.1016/s0165-0270(00)00335-6.
- [54] S. Kim, D. Putrino, S. Ghosh, and E. N. Brown, "A Granger causality measure for point process models of ensemble neural spiking activity," *PLoS Comput Biol*, vol. 7, no. 3, p. e1001110, Mar 2011, doi: 10.1371/journal.pcbi.1001110.
- [55] M. Sarmashghi, S. P. Jadhav, and U. Eden, "Efficient spline regression for neural spiking data," *PLoS One*, vol. 16, no. 10, p. e0258321, 2021, doi: 10.1371/journal.pone.0258321.
- [56] D. Mossman, "Resampling techniques in the analysis of non-binormal ROC data," *Medical decision making*, vol. 15, no. 4, pp. 358-366, 1995.
- [57] X. Robin *et al.*, "pROC: an open-source package for R and S+ to analyze and compare ROC curves," *BMC Bioinformatics*, vol. 12, p. 77, Mar 17 2011, doi: 10.1186/1471-2105-12-77.
- [58] B. C. Schwab *et al.*, "Neural activity during a simple reaching task in macaques is counter to gating and rebound in basal ganglia-thalamic communication," *PLoS Biol*, vol. 18, no. 10, p. e3000829, Oct 2020, doi: 10.1371/journal.pbio.3000829.
- [59] K. Bankiewicz, E. Oldfield, C. Chiueh, J. Doppman, D. Jacobowitz, and I. Kopin, "Hemiparkinsonism in monkeys after unilateral internal carotid artery infusion of 1-methyl-4-phenyl-1, 2, 3, 6-tetrahydropyridine (MPTP)," *Life sciences*, vol. 39, no. 1, pp. 7-16, 1986.
- [60] K. S. Bankiewicz, R. Sanchez-Pernaute, Y. Oiwa, M. Kohutnicka, A. Cummins, and J. Eberling, "Preclinical models of Parkinson's disease," *Curr Protoc Neurosci*, vol. Chapter 9, p. Unit9 4, May 2001, doi: 10.1002/0471142301.ns0904s09.
- [61] A. K. Engel and P. Fries, "Beta-band oscillations--signalling the status quo?," *Curr Opin Neurobiol*, vol. 20, no. 2, pp. 156-65, Apr 2010, doi: 10.1016/j.conb.2010.02.015.
- [62] T. Su and J. G. Dy, "In search of deterministic methods for initializing K-means and Gaussian mixture clustering," *Intelligent Data Analysis*, vol. 11, no. 4, pp. 319-338, 2007.
- [63] T. Donoghue *et al.*, "Parameterizing neural power spectra into periodic and aperiodic components," *Nat Neurosci*, vol. 23, no. 12, pp. 1655-1665, Dec 2020, doi: 10.1038/s41593-020-00744-x.
- [64] M. Deffains and H. Bergman, "Parkinsonism-related beta oscillations in the primate basal ganglia networks - Recent advances and clinical implications," *Parkinsonism Relat Disord*, vol. 59, pp. 2-8, Feb 2019, doi: 10.1016/j.parkreldis.2018.12.015.



- [65] A. Singh, "Oscillatory activity in the cortico-basal ganglia-thalamic neural circuits in Parkinson's disease," *Eur J Neurosci*, vol. 48, no. 8, pp. 2869-2878, Oct 2018, doi: 10.1111/ejn.13853.
- [66] J. Sarnthein and D. Jeanmonod, "High thalamocortical theta coherence in patients with Parkinson's disease," *J Neurosci*, vol. 27, no. 1, pp. 124-31, Jan 3 2007, doi: 10.1523/JNEUROSCI.2411-06.2007.
- [67] R. S. Eisinger *et al.*, "Parkinsonian Beta Dynamics during Rest and Movement in the Dorsal Pallidum and Subthalamic Nucleus," *J Neurosci*, vol. 40, no. 14, pp. 2859-2867, Apr 1 2020, doi: 10.1523/JNEUROSCI.2113-19.2020.
- [68] R. W. Rodieck, N. Y. Kiang, and G. L. Gerstein, "Some quantitative methods for the study of spontaneous activity of single neurons," *Biophys J*, vol. 2, no. 4, pp. 351-68, Jul 1962, doi: 10.1016/s0006-3495(62)86860-x.
- [69] V. Chan, P. A. Starr, and R. S. Turner, "Bursts and oscillations as independent properties of neural activity in the parkinsonian globus pallidus internus," *Neurobiology of disease*, vol. 41, no. 1, pp. 2-10, 2011.
- [70] J. A. Bilmes, "A gentle tutorial of the EM algorithm and its application to parameter estimation for Gaussian mixture and hidden Markov models," *International computer science institute*, vol. 4, no. 510, p. 126, 1998.
- [71] T. P. Mann, "Numerically stable hidden Markov model implementation," *An HMM scaling tutorial*, pp. 1-8, 2006.
- [72] D. Jurafsky and J. H. Martin, *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Pearson Prentice Hall, 2009.
- [73] I. H. Stevenson, "Omitted Variable Bias in GLMs of Neural Spiking Activity," *Neural Comput*, vol. 30, no. 12, pp. 3227-3258, Dec 2018, doi: 10.1162/neco\_a\_01138.
- [74] M. A. Kramer and C. J. Chu, "The 1/f-like behavior of neural field spectra are a natural consequence of noise driven brain dynamics," *bioRxiv*, p. 2023.03. 10.532077, 2023.
- [75] L. E. Wilson, J. da Silva Castanheira, and S. Baillet, "Time-resolved parameterization of aperiodic and periodic brain activity," *Elife*, vol. 11, Sep 12 2022, doi: 10.7554/eLife.77348.
- [76] C. Bauermeister, T. Schwalger, D. F. Russell, A. B. Neiman, and B. Lindner, "Characteristic effects of stochastic oscillatory forcing on neural firing: analytical theory and comparison to paddlefish electroreceptor data," *PLoS Comput Biol*, vol. 9, no. 8, p. e1003170, 2013, doi: 10.1371/journal.pcbi.1003170.
- [77] A. L. Hodgkin and A. F. Huxley, "A quantitative description of membrane current and its application to conduction and excitation in nerve," *J Physiol*, vol. 117, no. 4, pp. 500-44, Aug 1952, doi: 10.1113/jphysiol.1952.sp004764.
- [78] B. Ermentrout and A. Mahajan, "Simulating, analyzing, and animating dynamical systems: a guide to XPPAUT for researchers and students," *Appl. Mech. Rev.*, vol. 56, no. 4, pp. B53-B53, 2003.
- [79] D. J. Higham, "An algorithmic introduction to numerical simulation of stochastic differential equations," *SIAM review*, vol. 43, no. 3, pp. 525-546, 2001.
- [80] R. Rosenbaum *et al.*, "Axonal and synaptic failure suppress the transfer of firing rate oscillations, synchrony and information during high frequency deep brain stimulation," *Neurobiol Dis*, vol. 62, pp. 86-99, Feb 2014, doi: 10.1016/j.nbd.2013.09.006.

- [81] P. Dayan and L. F. Abbott, *Theoretical neuroscience: computational and mathematical modeling of neural systems*. MIT press, 2005.
- [82] S. Farhoodi and U. T. Eden, "The problem of perfect predictors in statistical spike train models," *Neurons, Behavior, Data analysis, and Theory*, vol. 5, no. 3, pp. 1-20, 2021.