A METHOD FOR AUTOMATED PERFORMANCE EVALUATION AND TUNING OF NEURAL OSCILLATION DETECTION ALGORITHMS

By

Charles Grimes Gerrity

Thesis

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

MASTER OF SCIENCE

in

ELECTRICAL ENGINEERING

December 18, 2021

Nashville, Tennessee

Approved:

Richard Alan Peters, Ph.D.

Thilo Womelsdorf, Ph.D.

ACKNOWLEDGMENTS

I would like to thank my advisors for helping me tremendously along the way. Professor Richard Alan Peters was ready to read through the edits that I sent him and help me when I struggled to write. Professor Thilo Womelsdorf was always eager to look through the results I was getting and making me think more critically about them. While not technically an advisor Christopher Thomas was invaluable in going through each of the questions I had along with helping me edit even when it was not easy.

I would like to thank all the lab mates I have had who have listened to my talks and presentations, giving valuable comments and making me think more clearly about this work.

I would also like to thank my friends and family for helping me along the way. Nick really helped me with my writing to get it into much better shape and I thank him tremendously for that. He was also there to support me through it all. I would also like to thank Connor for his abundant support throughout and listening to the ideas I had.

I would like to thank all of the other people who have helped who are too numerous to list here; your efforts were greatly appreciated

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

CAM Combined Angle Metric. vii, 58, 59, 64, 78, 79

CV Circular Variance. vii, 55–58, 64

ECoG Electrocorticography. 2

EEG Electroencephalography. 2

fMRI Functional Magnetic Resonance Imaging. 2

FN False Negative. vi, 27–30, 61

FP False Positive. vi, 27–30, 61, 64

LFP Local Field Potential. 2, 4, 5, 7

MD Mean Direction. vii, 55–58, 64

PPV Positive Predictive Value. 29

RMS Root Mean Square. iv, vii, ix, 33, 34, 36, 38, 40, 42, 44, 46, 48, 50, 52, 61, 63, 64

TP True Positive. vi, 27–30, 61, 64

TPR True Positive Rate. 30

GLOSSARY

Brownian noise Noise with spectral density proportional to $\frac{1}{f^2}$, where f represents frequency. 14, 15

Frequency Metric RMS of the instantaneous frequency residual. iv, vi, 33, 38–40, 62, 63, 78

Interpolated Relative Signal Metric RMS of the instantaneous relative signal residual with outliers interpolated. iv, vii, 34, 48–50, 62, 63

Magnitude Metric RMS of the instantaneous magnitude residual. iv, vi, 33, 36–38, 61, 63

plasticity Ability of a neuron to strengthen or weaken a synapse. 2

Relative Event Power to Bandpass Metric RMS of the signal residual relative to the RMS of bandpass signal. iv, vi, 33, 40–42, 62, 63

Relative Event Power to Wideband Metric RMS of the signal residual relative to the RMS of wideband signal. iv, vi, 33, 42–44, 62–64

Relative Frequency Metric RMS of the instantaneous relative frequency residual. iv, vii, 34, 52–54, 62, 63

Relative Magnitude Metric RMS of the instantaneous relative magnitude residual. iv, vii, 34, 50–52, 62, 63

Relative Signal Metric RMS of the instantaneous relative signal residual. iv, vi, 33, 44–46, 62, 63

Removed Relative Signal Metric RMS of the instantaneous relative signal residual with outliers removed. iv, vi, vii, 34, 46–48, 62, 63

Signal Metric RMS of the instantaneous signal residual. iv, vi, 33–36, 61–63

spike Neuronal spike firings indicating an action potential. 1, 2

spike sorting The process of using the shape of waveforms to determine individual neurons from background. 4

CHAPTER I

Introduction

The aim of this project is to detect and characterize electrical signals in the brain that represent local activity. These signals will be used to apply electrical stimulation to modulate activity and generate a behavioral response. The detection and characterization will inform the electrical stimulation because stimulation based on underlying neuronal dynamics has been shown to augment connections [1]. Augmenting fronto-striatal circuits may alter learning because these circuits are associated with learning either through action selection [2] or through the prediction error (the difference between the reward from a task and the predicted reward)[3].

I.1 Neurons

To describe neuronal signals, it is illustrative to start with the structure of the neuron itself. The typical neuron consists of three parts, the soma, the dendrites, and the axon [4]. The soma is the cell body where many of the organelles reside [4]. The dendrites and axons are the inputs and outputs of the neuron [4, 5]. Usually, an axon from one cell connects to a dendrite of another across a gap known as a synaptic cleft [4]. The connection across the gap is called a synapse [4]. The dendrites are extensions from the soma that act as antenna gathering input from other neurons[5]. The axon extends from the soma starting at a structure called the axon hillock and ending at the axon terminal which synapses to other neurons (or other cells) [4, 5].

The membrane of the neuron maintains a negative potential from the inside to the outside of the cell [4]. This is done through active and passive ion transport across the membrane [4, 6]. When the spike-initiation zone (the axon hillock in a typical neuron of the brain or spinal cord) is sufficiently depolarized by inputs from other neurons there is a rush of ions across the membrane [4, 5]. This electrical signal propagates along the axon, as an action potential or spike, releasing neurotransmitters into the synaptic cleft and taken up by the postsynaptic neuron [4, 5]. The neurotransmitter increases or decreases the membrane potential of the post-synaptic neuron [4, 5]. This either prevents or initiates another spike [4, 5]. This is the process of information transfer from the presynaptic neuron to the postsynaptic neuron [4, 5]. To summarize, one neuron is excited and sends it electrical excitement on to other neurons. This signal typically travels from axon to dendrites within a neuron and from dendrite to axons across neurons.

Neurons are a connection of inputs and outputs that form networks which process information. Some of the neurons in a network may fire synchronously [4]. This synchrony of spikes may occur either within a group of neurons following a central pacemaker or they may synchronize without a pacemaker to create a de facto timing function

that works to excite or inhibit each other [4]. The firing of these interconnected neurons generates electric potentials in the extracellular medium around the neurons, mostly from action potentials, but also from the aggregated ion diffusion which combine to form the Local Field Potential (LFP) [7].

I.2 Signal Measurement

The activity of neurons in the brain generates global signals that can be recorded via Functional Magnetic Resonance Imaging (fMRI), Electroencephalography (EEG), Electrocorticography (ECoG), and other methods [8]. More localized recording of neuronal activity is also done. The signal (spike train) from one neuron can be recorded by inserting a fine electrode into it [8]. The collective activity of a small network of neurons can be measured with electrodes that register the LFP [8]. This thesis focuses on measurements of the LFP of neural networks with electrodes.

The other types of recordings give other kinds of neural information. fMRI shows brain activity by measuring oxygen flow, which can indicate levels of activity based on energy usage [8]. The temporal resolution of fMRI is 1 to 2 seconds and the time between brain activity and blood oxygen response is 3 to 6 seconds [8]. EEG and ECoG are measures of the electric fields using a grid of electrodes either placed on the outside of the skull, EEG, or on top of the brain's surface, ECoG [8]. EEG represents the aggregate activity of several centimeters [9]. Penetrating electrodes are used to record single unit activity and the LFP [7, 8]. The LFP signal represents a local population of neurons and the signal can be different for electrodes separated by 1 mm [10] or about 250 μm [11]. ECoG is in between the spatial resolution of EEG and LFP recorded by penetrating electrodes [8].

I.3 Use of Measurements

Recorded signals, like LFP, are used to clarify the understanding of the brain based on the features found in them. One feature with a multitude of functions seen in various modalities is a burst event [12, 13]. Transient oscillatory activity, or burst events, has been defined in multiple ways. One definition is a short duration spontaneous increase in power of a particular frequency [14, 15]. These events can be found in a number of brain areas, including the primary somatosensory neocortex and the inferior frontal cortex [15], prefrontal cortex [12], striatum and motor-premotor cortex [16]. For this work the bursts are detected and characterized in the LFP measured from penetrating electrodes.

Direct electrical stimulation of the brain can alter brain function by either enhancing desired aspects or suppressing undesired aspects [1, 8, 17]. Because the LFP signal has oscillatory activity an instantaneous amplitude, frequency, and phase can be attributed to it [18]. Using the instantaneous phase, stimulation during burst events of LFP offer avenues to explore stimulation at certain phases, referred to as phase specific stimulation [1, 19–21]. Phase specific stimulation has been used to augment synaptic plasticity (the ability of a neuron to strengthen or

weaken a synapse) [1], enhance performance on spatial tasks relating to memory [22], reduce signal power in the beta band, 12.5 Hz to 30 Hz, alter reaction time [23], reduce oscillations in the beta band [19], and suppress tremors in Parkinson's patients [17]. Phase dependent stimulation had been shown to have an effect on information transfer in a simulated system as well [24]. Part of this project aims to improve the characterization of bursts for more accurate timing of phase specific stimulation.

Knowing when to stimulate, the next question is how to stimulate. There are many stimulation methods for transmitting encoded information to the brain such as transcranial magnetic stimulation, optogenetics, surface electrodes, etc. [8]. Depth electrodes are implanted into the brain tissue and typical stimulate with large currents or sustained high frequencies [25, 26]. This can cause unwanted byproducts to build up [27]. The electrode can be modeled as a capacitor and variable resistor in parallel with the brain tissue [27–31]. During stimulation, ions in the brain migrate to the electrode depending on the voltage until the capacitive component of the electrode is saturated [27–31]. After saturation, oxidation and reduction of the electrode and compounds in the brain will begin to occur, introducing potentially harmful chemicals into the brain [27]. While such stimulation is effective, it injects more energy than other methods that use the underlying brain signal, such as phase specific stimulation [1].

Phase specific stimulation requires the detection and characterization of burst events. There are many different methods for identifying bursts and characterizing them such as applying a threshold to the power in a specific frequency band [1, 32, 33]. Neuroscientists have performed experiments using phase specific stimulation to test their hypotheses [1, 32, 33]. The strength of these results depends on how well the bursts are characterized and detected [34]. The parameters of a method are typically selected ad hoc [1, 34, 35]. Validation of methods employed is typically not the aim of these kinds of experiments, which may not reveal all downsides of the methods [1, 32, 33]. An example of one group's validation data has been a chirp over a narrow band [36]. A chirp or sweep signal is a signal that has an increase or decrease in frequency over the time of the signal [37]. This is a typical method but may be too simple for our purposes because it does not entail a sharp increase in amplitude over a brief time similar to a burst event. Current comparison of methods using performance metrics have significant limitations. Some comparisons do not separate detection and characterization [33, 38]. Others use a narrow range of metrics for comparison [39], or only compare a small number of methods [40, 41]. It is important to analyze both detection and characterization which can be used for burst classification and also provide valuable characteristics such as phase that should be accurate for the stimulation methods. This project will include a range of metrics on validation data that more closely resembles recorded data in addition to using recorded data.

Because of the above problems with evaluating performance, we want a systematic and broadly-applicable way of comparing the performance of burst detection and characterization algorithms. An important goal of the work

reported in this thesis is a framework for comparing and tuning LFP burst detection and characterization algorithms. Similar approaches have been proposed for other problems. Frameworks for another neuroscience signal processing component, spike sorting, has been developed by other researchers [42]. Our main goal is to determine the highest performing algorithm for use in phase specific stimulation experiments through a robust way to measure algorithm performance and tune existing algorithms with future work leading to the development of new extensions of these algorithms. The performance before tuning will come from literature implementations. The comparison will be done by extending an existing framework for burst detection and characterization methods developed in our lab. The extensions will include more methods, and addition of metrics for comparison and tuning. A number of performance metric values will facilitate comparison. A suitable metric will be chosen from these to use as an objective function for tuning the parameters of the method eliminating the need to hand-tune method specific parameters.

CHAPTER II

Background

In this work five methods are explored for LFP burst detection and characterization. These methods represent different signal processing approaches. The methods will be called in this work Hilbert Magnitude, Peak and Trough, Wavelet, Template Fitting, and Frequency Stability. Frequency Stability method is an unpublished work by Womelsdorf lab [43], using their library functions [44].The background for understanding these methods is explained in the following sections. The implementation of each method is explained in the following chapter.

This section illustrates how the methods are applied to a signal and how to evaluate the method. The methods are categorized as having two parts as shown in Equation II.1. The j-th method, $M_j(\cdot;\tau_j)$ is the composition of the j-th characterization,$C_j(\cdot;\tau_j)$ and j-th detection,$D_j(\cdot;\tau_j)$ functions. These functions have tuning parameters, which are aggregated into the vector τ_j, where some of the parameters adjust the detection and some adjust the characterization. The method acts on a signal, $S(t)$, to produced a detected and characterized signal, $M_j(t)$.

$$M_j(t) = M_j(S(t);\tau_j) = C_j(D_j(S(t);\tau_j);\tau_j) = (C_j \circ D_j)(S(t);\tau_j) \tag{II.1}$$

The first part is a detection method identifying where the burst is as shown in Equation II.2. The detection function outputs when there is a burst, which is shown here as a binary function where 1 indicates the detection of a burst and 0 does not indicate a burst.

$$D_j(t) = D_j(S(t);\tau_j) = \begin{cases} 1 & \text{if burst detected by } D_j(\cdot;\tau_j) \\ 0 & \text{else} \end{cases} \tag{II.2}$$

The second is a characterization method, which identifies the instantaneous amplitude $A(t)$, instantaneous phase $\phi(t)$, and instantaneous frequency $f(t)$ of the burst as shown in Equation II.3. The characterization function has a domain only over the time points where the detection function is 1 otherwise there is no need for characterization since no burst is detected.

$$C_j(t) = C(D_j(S(t);\tau_j);\tau_j) = \begin{cases} A(t)\sin(\phi(t)) \text{where } \phi(t) = \int_0^t 2\pi f(x)\,dx + \theta & \text{if } D_j(t) = 1 \\ \text{Undefined} & \text{else} \end{cases} \tag{II.3}$$

In implementations of these methods, the tuning parameters are usually set ad-hoc without being tuned first. The parameters will alter how the bursts are detected and characterized.

Once a method is applied to the signal the quality of the method is assessed using performance metrics, μ. Equation II.4 shows that the application of the i-th metric function, $\mu^i(\cdot)$, on a detected and characterized signal using the j-th method, $M_j(t)$, yields the scalar value for the i-th metric on the j-th method, μ_j^i.

$$\mu_j^i = \mu^i(M_j(t)) \tag{II.4}$$

All of the metric functions are applied to the result of the j-th method leading to the set of metric values for the j-th method as shown in Equation II.5. This set of metric values for the method are used for comparison to other methods. Single metric values or combinations of metric values can also be used to optimize the tuning parameters of a method by comparing the results of a method to itself with different values for tuning parameters.

$$\mu_j = \{\mu_j^1, \ldots, \mu_j^i, \ldots, \mu_j^n\} \tag{II.5}$$

Figure II.1 shows how the methods are evaluated. The burst is detected and characterized by the chosen method. Then this method is evaluated using the chosen metric. This process is performed for all methods $j \in \{1, \ldots, L\}$ and all metrics $i \in \{1, \ldots, K\}$.

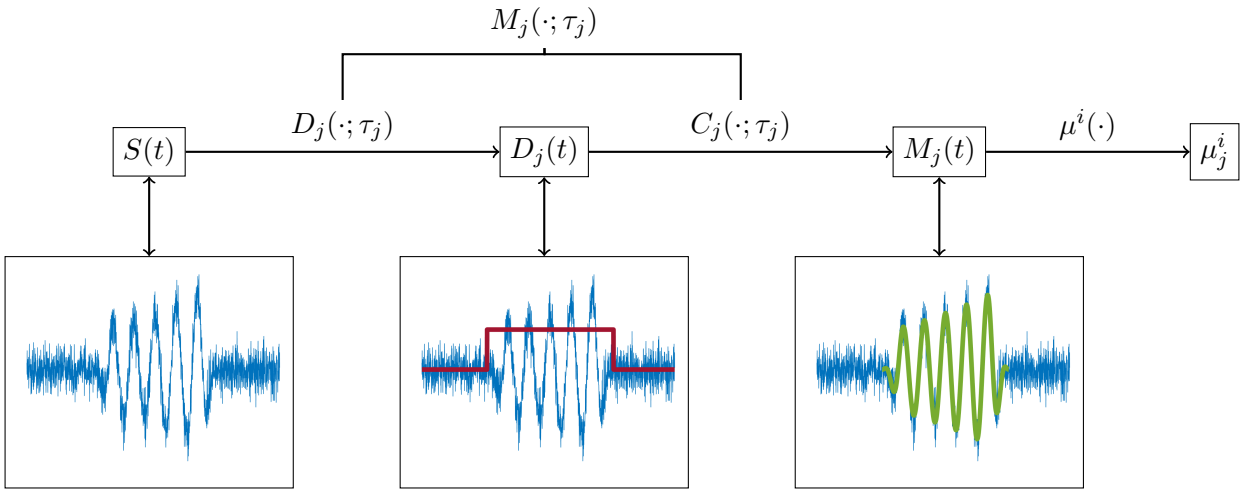

Figure II.1: This diagram shows the process of the metric experiment. The method, M_j, is applied to the signal, $S(t)$. M_j is composed of two parts, the detection function, D_j, and the characterization function, C_j. These components use the tuning parameters, τ_j, which are associated with the method. These are applied to $S(t)$ to first get the detected events, represented by $D_j(t)$ then the events are characterized to get the detected and characterized signal by the j-th method, $M_j(t)$. The i-th metric function μ_i is applied to the detected and characterized events of the j-th method to obtain the metric value, μ_j^i.

To compare and validate burst detection and characterization methods metrics will be used (e.g. false positive rate used in spike train burst detection [45], phase jitter [20, 34], confusion matrix statistics [46, 47]). In addition to comparing and validating a method metrics can be used to tune the parameters of a method (e.g F_β [33]). Tuning optimization is often handpicked using an acceptable value for a metric, but more systematic approaches offer a way to tune multiple parameters at the same time (e.g.grid search method described in [33] or genetic algorithm optimization described in [20]).

Figure II.2 shows how the methods will be tuned. A metric is chosen beforehand, $\mu^*(\cdot)$ that will be used as the objective function for tuning. The calculation of the metric proceeds the same as above. The difference is that the selection of τ_j is determined by an optimizing function $\Omega^k)(\cdot)$. This process is repeated until the criteria for optimization are satisfied leading to τ_j^*, which represents the tuning parameters for the j-th method that are optimal according to the selected metric.

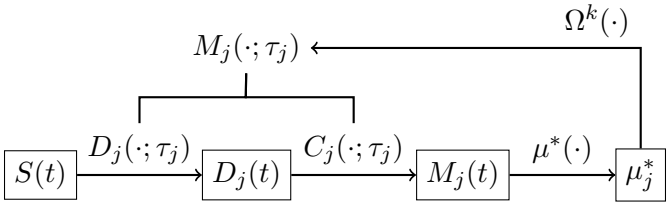

Figure II.2: This diagram shows the process of the tuning experiment. The method, M_j, is applied to the signal, $S(t)$, similarly to the metric experiment. M_j is composed of two parts the detection function, D_j, and the characterization function, C_j. These components use the tuning parameters, τ_j, which are associated with the method. These are applied to $S(t)$ to first get the detected events, represented by $D_j(t)$ then the events are characterized to get the detected and characterized signal by the j-th method, $M_j(t)$. The difference between the tuning and the metric experiment is that a metric has been chosen from the set μ^i to be used as the objective function for tuning, which is labeled, μ^* here. The value of μ^* is optimized using an optimization function, Ω^k, over its input τ_j. This alters the tuning parameters for method M_j, leading to different detection and characterization by C_j and D_j. This in turn changes the value of μ^*. This process is repeated until the value of μ^* has reached a stopping criteria, which outputs τ_j^*, the optimal set of tuning parameters for the j-th method using μ^* as the objective function.

II.1 Detection and Characterization Approaches

The Hilbert Magnitude method is fairly common in LFP burst detection that uses the Hilbert transform to detect and characterize the burst[48]. The Peak and Trough method uses the waveform in the time domain to characterize the signal and then detect the characterized signal[23, 33, 36, 48]. The Wavelet method uses wavelets to generate a time-frequency representation of the signal for both detection and characterization [35, 39]. The Template Fitting method uses a Hilbert magnitude method for detection, with template fitting for characterization[1]. The Frequency Stability method detects the stability of the Hilbert frequency and uses the Hilbert transform for characterization [43].

In the following sections the background information for these burst detection and characterization methods is explored.

II.1.1 Hilbert Magnitude

An analytic signal is a signal that has no negative frequency components. The analytical signal is useful because it gives the in-phase and quadrature components, which can give the instantaneous magnitude and phase of oscillating signals. Because real signals inherently have symmetric frequency components then an analytic signal can be constructed from the Fourier transform of a real signal, $s(t)$, based on Equation II.6 and Equation II.7. The Fourier transform of the analytic signal is described in Equation II.6, which is a result of the Fourier transform of a real signal, $s(t)$, having only real components. The analytic signal, $s_a(t)$, is the inverse Fourier transform of, $S_a(f)$. Equation II.7 shows this algebra, which demonstrates that the real part of the analytic signal is the original signal, $s(t)$, but the imaginary part is exactly the Hilbert transform of the original signal, $\hat{s}(t)$.

With a given input signal the Hilbert transform is used to generate an analytical signal, which is represented as real (in-phase) and imaginary (quadrature) components. From that, the instantaneous magnitude is taken as the norm and the instantaneous phase is taken as the arctan of the real and imaginary components. The frequency is calculated as the derivative of the unwrapped phase. The unwrapped phase is smoothed in order to reduce noise[48].

$$S_a(f) \triangleq \begin{cases} 2S(f), & \text{for } f > 0 \\ S(f), & \text{for } f = 0 \\ 0 & \text{for } f < 0 \end{cases}$$
$$= \underbrace{2u(f)}_{1+\text{sgn}(f)} S(f) = S(f) + \text{sgn}(f)S(f) \tag{II.6}$$

$$
\begin{aligned}
s_a(t) \triangleq & \mathcal{F}^{-1}[S_a(f)] \\
= & \mathcal{F}^{-1}[S(f) + \text{sgn}(f)S(f)] \\
= & \underbrace{\mathcal{F}^{-1}\{S(f)\}}_{s(t)} + \overbrace{\underbrace{\mathcal{F}^{-1}\{\text{sgn}(f)\}}_{j\frac{1}{\pi t}} * \underbrace{\mathcal{F}^{-1}\{S(f)\}}_{s(t)}}^{\text{convolution}} \\
= & s(t) + j \underbrace{\left[\frac{1}{\pi t} * s(t)\right]}_{\mathcal{H}[s(t)]} \\
= & s(t) + j\hat{s}(t)
\end{aligned} \tag{II.7}
$$

The analytic signal has properties such that the instantaneous amplitude, frequency, and phase can be calculated. The instantaneous amplitude is calculated per Equation II.8, the instantaneous phase is calculated per Equation II.9, and the instantaneous frequency is calculated per Equation II.10 [49, 50].

$$A(t) = |s_a(t)| = \sqrt{s(t)^2 + \hat{s}(t)^2} \tag{II.8}$$

$$\phi(t) = \angle s_a(t) = \arctan\left(\frac{\hat{s}(t)}{s(t)}\right) \tag{II.9}$$

$$f(t) = \frac{d\phi(t)}{dt} \tag{II.10}$$

For the Hilbert Magnitude detection, a burst is detected when the magnitude of the signal obtained from the Hilbert transform surpasses a threshold. A multiple of the standard deviation of magnitude [1], standard deviation of power [23], or percentile of amplitude[39, 51] are used to determine thresholds.

II.1.2 Peak and Trough

Peak and Trough methods are described in [23, 33, 36, 48]. This method aims to bypass the use of sinusoidal basis functions that may alter the estimates of the envelope and carrier functions if the burst is not accurately described by a sinusoid. Peak and trough and zero-crossing based detection estimates signal magnitude, frequency, and phase without computing the analytic signal[48]. Some implementations use midpoints with the aid of zero-crossings[33] and another method using zero-crossings of the derivative of the smoothed signal to find peaks and troughs [36]. An example signal is shown in Figure II.3.

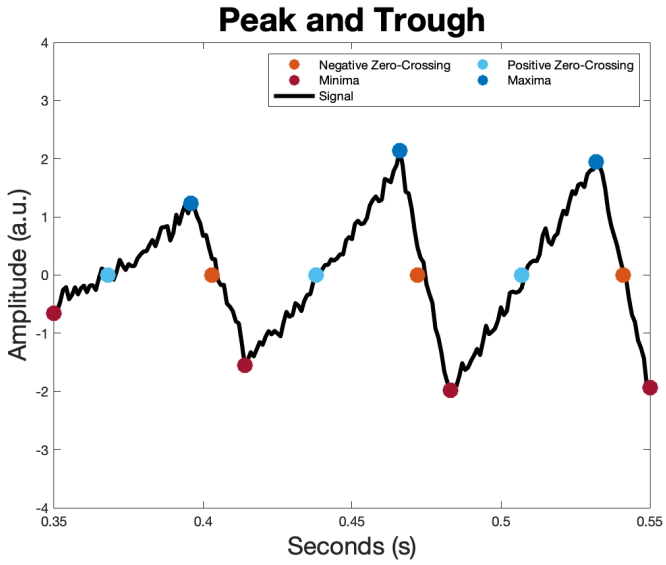

Figure II.3: This figure shows an example of the Peak and Trough method. Here the black saw tooth line represents the signal to analyze. Zero crossings of the signal are marked and noted as either positive (positive slope) or negative (negative slope). These values divide the signal into distinct cycles. The maximum or minimum of the signal is marked depending on the value of the prior zero crossing, maximum for a positive zero crossing and minimum for a negative zero crossing.

With these methods the phase is assigned based on the point and interpolated between points [33]. The magnitude is taken as a function of the amplitude in the neighborhood of the point [33]. The frequency can be derived from the period, which comes from the time between the designated points [33].

II.1.3 Wavelet

An example wavelet method is described in [14, 35]. A wavelet transform is used to obtain the spectrogram of the signal. This shows the evolution of the signal power in time and frequency. After obtaining this two dimensional representation of the signal various templates can be fit onto it to isolate components in the signal, features can be extracted, or thresholds applied [46]. This will result in the detection of events in the signal that are determined with both time and frequency.

Ψ is the initial wavelet, the transform then dilates and shifts that template to get other wavelets that are convolved with the signal to get a spectrogram, a time-frequency representation [52]. The wavelet transform, $W(u, s)$, of the signal, $x(t)$, is shown in Equation II.12. The wavelet can be any signal that is concentrated in time with zero mean and for use in the transform is normalized to 1 [52, 53]. The dilation and translation of the "Mother" wavelet are done through parameters s and u as shown in Equation II.11 [52, 54]. To convert to a time-frequency response the dilation and translation parameters are varied with the translation representing the time and the dilation being

inversely proportional to frequency [52]. Because of the dilation parameter the support size of the transform scales with the frequency. At lower frequencies, larger values of s, the wavelet is expanded and the time support of the wavelet is larger, but at higher frequencies, smaller values of s, the wavelet is contracted and the time support of the wavelet is smaller. This gives the advantage of having resolution that adapts to the given frequency, a component that is not found in other time-frequency methods such as the short time Fourier transform that uses a fixed window size [52, 54].

$$\psi_{u,s}(t) = \frac{1}{\sqrt{s}}\psi(\frac{t-u}{s}) \tag{II.11}$$

$$W(u,s) = \int \frac{1}{\sqrt{s}}\psi^*(\frac{t-u}{s})x(t)dt \tag{II.12}$$

Many functions can be used for the "Mother" wavelet, but here a handful of examples will be shown. The Gabor wavelet is shown in Equation II.13 [53]. the Morlet wavelet is shown in Equation II.14 [53]. The Morse wavelet is shown in Equation II.15 and is defined in the frequency domain [55]. In this equation $U(\omega)$ is the Heaviside step function and $\alpha(\beta,\gamma)$ is defined with Equation II.16 .

$$\psi(t,\sigma,\eta) = \frac{1}{\sqrt{\pi\sigma}}e^{-\frac{t^2}{2\sigma^2}}e^{i\eta t} \tag{II.13}$$

$$\psi(t,\eta) = e^{-\frac{t^2}{2}}e^{i\eta t} \tag{II.14}$$

$$\Psi(\omega,\beta,\gamma) = U(\omega)\alpha(\beta,\gamma)\omega^{\beta}e^{\omega^{\gamma}} \tag{II.15}$$

$$\alpha(\beta,\gamma) = 2\left(\frac{e\gamma}{\beta}\right)^{\frac{\beta}{\gamma}} \tag{II.16}$$

The wavelet transform using the Morse wavelet provides the most accurate instantaneous phase because the Morse wavelet is an exactly analytic wavelet [55]. The other wavelets are close but not exactly analytic and will have negative frequency components leading to a mischaracterization of instantaneous frequency and phase [55].

II.1.4 Template Fitting

Template fitting methods use a predefined shape and match it to the input signal. Templates can be either parametric or non-parametric. The non-parametric templates represent an extension of dictionary based offline algorithms.

These methods use parts of the signal as templates to match with other parts of the signal. The signal templates can be generated in a variety of ways including principal component analysis [56] and signal statistics [57].

Other template fitting methods use a parametric approach to template fitting. One method uses a Hilbert magnitude based detection and characterizes by fitting a sine curve as shown in Equation II.17 to the signal [1]. This is similar to dictionary methods that use a basis function with free parameters to fit the detected burst. In the example template in Equation II.17, A, ϕ_0, and T represent the free parameters.

$$y(t; A, \phi_0, T) = A \sin(\phi_0 + 2\pi \frac{t}{T}) \tag{II.17}$$

To fit this template to the signal, $x(t)$, the difference between the template and the signal is taken as the error. The error is squared and summed and shown in Equation II.18. The detector determines the start and end time for the burst, which correspond with i and M in this equation.

$$SSE(A, \phi_0, T) = \sum_{i=1}^{M} (x(i) - y(i; A, \phi_0, T))^2 \tag{II.18}$$

An optimization is performed over the free parameters to minimize the sum of squared errors. These free parameters are then used to characterize the magnitude, frequency, and phase of the burst.

II.1.5 Frequency Stability

Frequency stability and phase stability are methods that detect bursts when the phase or frequency is roughly constant. First these methods extract the instantaneous frequency or phase of the signal. Then they transform this signal to get one that is constant when a distinct burst is occurring.

In one method the instantaneous phase is calculated from the continuous Gabor transform of the signal [58–60]. This transform is similar to the wavelet transform from Equation II.12 using a Gabor wavelet, Equation II.13, for the mother wavelet. After the insantaneous phase is obtain the rotated phase is caalculated using Equation II.19. The rotated phase, $\phi_R(t, \omega)$, represents the time shifted phase relative to a starting phase. Δt represents the time from the reference phase and $T = 1/\omega$ [58–60].

$$\phi_R(t, \omega) = \phi(t, \omega) - \Delta\phi = \phi(t, \omega) - 2\pi \bmod(\Delta t, T)/T \tag{II.19}$$

The rotated phase will stay within a small range when the signal is coherent indicating that there is a burst. To determine a burst the rotated phase must stay within a certain range such as 45° [58–60].

Another method developed in our lab uses frequency stability to detect bursts [44]. For the Frequency Stability

detection, noise is added to the signal and then the frequency is obtained from the Hilbert transform just as with the Hilbert Magnitude method. The frequency is only stable when the signal is coherent. The addition of noise reduces strong correlation between neighboring samples. This allows for the detection of bursts when the estimate of the frequency stabilizes. This method was selected because it has a large number of tuning parameters that lends itself to the analysis.

II.2 Optimization Approaches

There are many approaches to solving optimization problems. The methods for solving these include approaches based on the local topography such as gradient based methods, random based methods such as random search, and brute force methods such as grid search[61–63]. There are many more approaches, but here two methods for optimization are used. The first is a grid search method and the second is a "creeping" random search.

Grid search methods first require a search space to be determined a priori if the domain is infinite for the function to optimize [61]. Once a suitable domain for all of the parameters is chosen then a particular permutation of the parameters is chosen, which will determine the order that they are optimized for univariate optimization[64]. After this the each domain is split into probe points and the combination is tested for one parameter, univariate, to determine the best combination of inputs[61, 64]. The univariate grid search prevents the combinatorial explosion of having to search over higher dimensional domains[61].

"Creeping" random search is a type of random search that looks for optimal points in an area around the current optimal point[65]. The initial point is given and then a random vector is drawn from a given distribution and added onto the initial point[66]. The objective function is evaluated at this new point and if it is higher (lower) for maximization (minimization) than the previous point, the new point is accepted and the next random vector is added to this point otherwise a new random vector is added to the initial point[67–69]. The probability density function of the random vector is different in various applications and for some the random vector will become smaller and smaller as the optimization proceeds similar to simulated annealing[61–63].

CHAPTER III

Implementation of Methods

The goal of the project is to implement several burst detection and characterization methods, measure their performance, and automatically tune their performance in a useful way. Tuning and performance measurement are described in subsequent chapters. This chapter describes the methods that were implemented and the datasets that were tested. This includes the synthetic and real data used for the analysis, the model of the data and methods, and each of the previously described methods.

The methods used in this work Hilbert Magnitude, Hilbert Frequency, Peak and Trough, Cosine Template, and Wavelet.

The datasets are one synthetic dataset with realistic noise spectrum and events, and one real primate dataset, described in this chapter.

III.1 Description of Data

III.1.1 Signal Model

There are many ways to model the signal. Some have used Gabor atoms, which are sinusoids of the desired frequency modulated by a Gaussian [54, 70]. For this work, it is convenient to model the signal as having each event represented as a wavelet-like event with additive red noise, also known as Brownian noise. The event has a representative envelope function and carrier function. The carrier function has an associated frequency function and phase while the envelope function has an associated magnitude function, which is nonzero only within a finite interval . The magnitude and frequency functions are not necessarily constant over the event. The events occur at various time points in the overall signal and may be overlapping.

Red noise is defined as having a power spectrum that falls off with the inverse of the square of frequency.

$$S_A(t) = \left(\sum_{n=1}^{N} \Psi_n(A_n(t), f_n(t), \phi_{0n}, t_{0n}) \right) + \epsilon(t) \tag{III.1}$$

$$S_\epsilon(f) \propto \frac{1}{f^2} \tag{III.2}$$

In Equation III.1 the recorded local field potential is represented as $S_A(t)$. The event function Ψ has input values describing the envelope $A_n(t), t_{0n}$ and inputs describing the carrier function $f_n(t), \phi_{0n}$. In this representation there are N bursts which are summed to produce the recorded signal along with a background red noise error term, $\epsilon(t)$,

that has a power spectrum described by Equation III.2. This model does not specify the type of envelope or its characteristics, which are assumed in the methods used to estimate the events. The carrier function is assumed to be periodic, from the parameters that it has a frequency and phase, but the frequency is not fixed over the burst event. The various methods will assume functional forms for the parameters, reducing the generality.

III.1.2 Synthetic Data

The synthetic signal is generated by the existing framework developed by the Womelsdorf lab [44]. Various parameters for the synthetic signal are first defined. The next step is to generate the background red noise of the signal, which is done using the function dsp.ColoredNoise() from MATLAB. The signal is then band pass filtered over the wide band (1 to 40% of the sampling frequency).

The next step is to add bursts to the background noise trace. Bursts are added according to the specified parameters. The magnitude envelope is a cosine roll off and roll on with either a logarithmic or linear ramp between the starting and ending magnitude values. The frequency is also a logarithmic or linear ramp between the starting and ending values. The phase is determined from the frequency and a phase offset term. The phase is used to generate the carrier function, which is the cosine of the phase. The synthetic data is illustrated in Equation III.3 where the Amplitude, $A_k(t)$, is non-zero for a finite interval only during the burst event.

Each burst is recorded in an event list, which includes the time of the burst, magnitude, frequency, phase and the parameters used to generate each burst. This event list is the ground truth for the synthetic data [44].

$$S(t) = \left(\sum_{k=1}^{K} \Psi_n(A_k(t), f_k(t), \phi_{0k}, t_{0k}) \right) + \epsilon(t) \tag{III.3}$$

III.1.3 Real Data - Non-Human Primate

The real dataset comes from recordings in non-human primates. The recording is segmented by individual trials, lasting less than 10 seconds, that occurred during the recording experiment. If an event is detected near the edge of a segment it is not counted as it may have been an artifact.

To show that the synthetic and real datasets are similar Figure III.1 shows the power spectrum of two signals from the datasets. The real data appears to show a steeper drop off with higher frequencies that the synthetic data indicating that the noise may drop off more than $\frac{1}{f^2}$ as is the case with Brownian noise. This shows that the synthetic dataset has some similar spectral qualities to the real dataset.

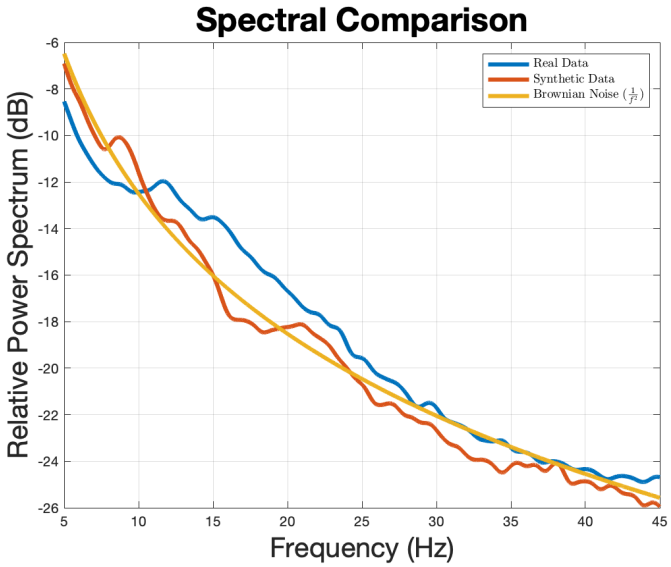

Figure III.1: This figure shows the power spectrum of the real and the synthetic data. The frequency range stops at 50 Hz to because higher frequencies are not discussed in this work. The two spectra align well with the real data shower a slightly steeper drop off for increased frequency. The real data also shows an increase in power between 10 and 20 Hz which aligns with the band of interest in this work

III.1.4 Method Model

Various methods attempt to detect the events, generating their own event lists. Each location is then characterized to obtain the magnitude, frequency, and phase estimates. The methods have multiple tuning parameters τ_j that alter the events detected and how they are characterized.

$$M_j(S(t), \tau_j) = M_j(t) = \left(\sum_{l=1}^{L} \Psi_l(A_l(t), f_l(t), \phi_{0l}, t_{0l}) \right) \tag{III.4}$$

The method model, illustrated in Equation III.4, shows that for a given method, each burst is recorded in an event list, which includes the time of the burst, magnitude, frequency, phase and the parameters used to generate each burst. L bursts are detected by the method, which extracted events from the signal, $S(t)$. Each method has a different way of modeling the events, which is shown through how it uses the parameters to construct the burst in the $\Psi(\cdot)$ function.

III.2 Burst Methods

Most of these work on signals that have been band-pass-filtered. They do a detection step to recognize when oscillations are occurring in-band and then a parameter extraction step to get magnitude, frequency, and phase during the event. The wavelet method applies to the wide band, but analyzes events on a per band basis just the same.

III.2.1 Hilbert Magnitude

Analytic magnitude and phase are taken as being correct obtained from the Hilbert transform of the bandpass filtered signal. This is typically done in a standard band such as the beta band, 12.5 Hz to 30 Hz [12, 23]. Frequency is calculated as the derivative of smoothed phase. Magnitude-based thresholding is done similar to [48]. This method will be referred to as the "Hilbert Magnitude" algorithm.

To start the Hilbert Magnitude method the signal is bandpass filtered to the desired band. The Hilbert transform is applied to generate the analytic signal. The instantaneous magnitude is obtained from the magnitude of the analytic signal. The instantaneous phase is obtained similarly as the phase of the analytic signal. The instantaneous frequency is found as the derivative of the instantaneous phase with smoothing applied. A base magnitude is obtained from lowpass filtering the magnitude signal. This will give a reference to determine any excursions that would indicate a burst. The power of the instantaneous amplitude and the lowpass instantaneous amplitude are obtained. The ratio between these two determines whether a burst is detected if the ratio surpasses a value set by the dB_{peak} tuning parameter. If a burst is detected its end points are extended to a lower threshold of this ratio set by the dB_{end} tuning parameter. To remove spurious events and combine small dropouts that should be from one event, gaps between detected events that are larger than $maxdrop$ are filled in and events shorter than $maxglitch$ are removed. This gives a list of events by when they start and how long they occur.

A further extension of Hilbert magnitude and Hilbert frequency uses template fitting to get event parameters. This is the implementation provided by the `wlBurst_v2` library used for experiments. Template fit used is "chirp ramp" [44]. Unless otherwise noted, Hilbert Magnitude and Hilbert Frequency used for experiments have "chirp ramp" parameter fits for detected events.

For characterization of the detected events, the waveform is parameterized using starting and ending amplitude, starting and ending frequency, staring phase, roll on, roll off, sample start and duration. These parameters are used to generate waveforms that are compared to the bandpass signal. To calculate these parameters a grid search is performed. Using the values from that step creeping random search is performed to obtain better estimates.

Algorithm 1 Hilbert Magnitude Method

1: Signal is bandpass filtered for the specified band, generating $S_{bp}(t)$
2: The analytic signal, $s_a(t)$ is obtained using the Hilbert transform
3: $A(t) = abs(S_a(t))$
4: $\theta(t) = angle(S_a(t))$
5: $f(t) = \frac{d\theta(t)}{dt}$
6: $A_{slow}(t)$ comes from $A(t)$ lowpass filtered with time constant τ_{DC}
7: $P_{slow}(t) = A_{slow}(t)^2$
8: $P_{fast}(t) = A(t)^2$
9: The ratio of $P_{fast}(t)$ to $P_{slow}(t)$ with a putative burst detected if the ratio exceeds the threshold dB_{peak}
10: Extend detected bursts by comparing to another threshold dB_{end} only extended previous bursts to this lower threshold
11: Combine bursts separated by time gap smaller than $maxdrop$
12: Remove events shorter than $maxglitch$
13: This gives a list of M burst events
14: **for** each event, i, of M bursts **do**
15: Do Grid Search for initial amplitude, frequency, phase, roll on, roll off, sample start and duration
16: Do Creeping Random Search for starting and ending amplitude, roll on, roll off, sample start and duration
17: Do Creeping Random Search for starting and ending amplitude, starting and ending frequency, staring phase, roll on, roll off, sample start and duration
18: **end for**
19: **return** List of M events with A_i, f_i, θ_i for each event

III.2.2 Hilbert Frequency

This is an extension of "Hilbert magnitude" algorithm that uses local variance of the frequency estimate for detection rather than magnitude. During clean oscillations, variance is low, but outside of oscillations, variance is high. This is an unpublished work by Womelsdorf lab [43], using their library functions [44]. This method will be referred to as the "Hilbert Frequency" algorithm.

To start the Hilbert Frequency method the signal is bandpass filtered to the desired band. The Hilbert transform is applied to generate the analytic signal. The instantaneous magnitude is obtained from the magnitude of the analytic signal. The instantaneous phase is obtained similarly as the phase of the analytic signal. The instantaneous frequency is found as the derivative of the instantaneous phase with smoothing applied. A new signal is generated by adding white noise in a band determined by the tuning parameters Lim_{lower} and Lim_{upper} with power determined by SNR_{noise}. The instantaneous frequency of this new signal is obtained in the same way as described earlier. A slowly changing instantaneous frequency of this noisy frequency is obtained by lowpass filtering with corner frequency determined by the parameter, Q_{long}. Next variance as a function of time is calculated. A slow and fast changing variance are calculated using corner frequencies set by Q_{short} and Q_{long} respectively. The ratio of the slow to fast variance are compared to a threshold determined by dB_{peak}. f a burst is detected its end points are extended to a lower threshold of this ratio set by the dB_{end} tuning parameter. To remove spurious events and combine small

dropouts that should be from one event, gaps between detected events that are larger than $maxdrop$ are filled in and events shorter than $maxglitch$ are removed. This gives a list of events by when they start and how long they occur. These events are characterized in the same manner as with Hilbert Magnitude (Section III.2.1).

Algorithm 2 Hilbert Frequency Method

1: Signal is bandpass filtered for the specified band, generating $S_{bp}(t)$
2: The analytic signal, $s_a(t)$ is obtained using the Hilbert transform
3: $A(t) = abs(S_a(t))$
4: $\theta(t) = angle(S_a(t))$
5: $f(t) = \frac{d\theta(t)}{dt}$
6: White noise is added to $S_{bp}(t)$ generating $S_n(t)$ in the band $[Lim_{lower}, Lim_{upper}]$ with power such that the SNR is SNR_{noise}
7: The instantaneous frequency $f_n(t)$ is obtained similarly to the above using $s_n(t)$
8: Calculate a slowly changing mean $f_{n,\mu}(t)$ by low pass filtering $f_n(t)$ with corner frequency based on Q_{long}
9: $Var_n(t) = (f_n(t) - f_{n,\mu}(t))^2$
10: Obtain $Var_{n,fast}$ by low pass filtering $Var_n(t)$ with corner frequency based on Q_{short}
11: Obtain $Var_{n,slow}$ by low pass filtering $Var_n(t)$ with corner frequency based on Q_{long}
12: The ratio of $Var_{n,slow}$ to $Var_{n,fast}$ with a putative burst detected if the ratio exceeds the threshold dB_{peak}
13: Extend detected bursts by comparing to another threshold dB_{end} only extended previous bursts to this lower threshold
14: Combine bursts separated by time gap smaller than $maxdrop$
15: Remove events shorter than $maxglitch$
16: This gives a list of M burst events
17: **for** each event, i, of M bursts **do**
18: Do Grid Search for initial amplitude, frequency, phase, roll on, roll off, sample start and duration
19: Do Creeping Random Search for starting and ending amplitude, roll on, roll off, sample start and duration
20: Do Creeping Random Search for starting and ending amplitude, starting and ending frequency, staring phase, roll on, roll off, sample start and duration
21: **end for**
22: **return** List of M events with A_i, f_i, θ_i for each event

III.2.3 Peak and Trough

This work uses an implementation of published work[33]. This method will be referred to as the "peak and trough algorithm".

To obtain estimates of the magnitude, frequency, and phase the method illustrated in Figure III.2 must locate various marker points in the signal. The first step is to band pass filter the signal in the band of interest. This allows for the zero-crossings to be calculated. The next steps strictly use the wide band signal. The extrema are determined between each zero-crossing, maxima or minima depending on the type of zero-crossings. The next step is to get the flank points, which are the points at which the signal reaches halfway between adjacent extrema. This gives a value that will be different from the zero-crossing the less sinusoidal the input waveform is.

These points are used to determine magnitude, phase, and frequency. The magnitude is half the absolute differ-

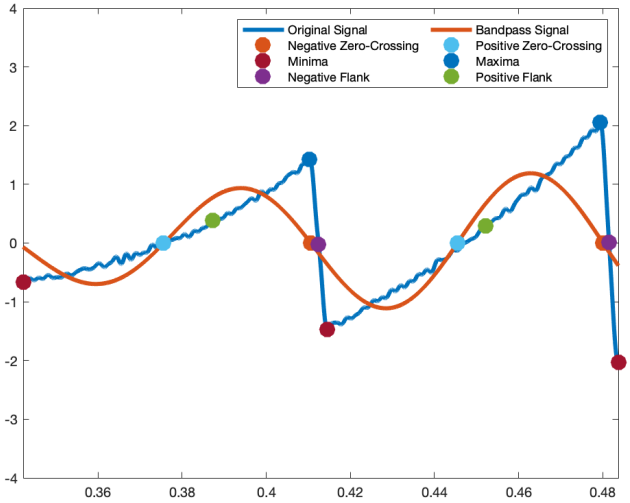

Figure III.2: This figure shows an example of the peak and trough method described in [33]. Here the blue saw tooth line represents the signal to be analyzed and the orange line is the band passed version of the same signal in the band of interest. Zero crossings of the band pass signal are marked and noted as either positive (positive slope) or negative (negative slope). These values divide the wide band signal into distinct cycles. The maximum or minimum of the wide band signal is marked depending on the value of the prior zero crossing, maximum for a positive zero crossing and minimum for a negative zero crossing. The flank points represent the signal value that is the average of the adjacent extrema. Because the original signal is not always symmetric the flank point will not always correspond with the zero crossing. As the original signal becomes more asymmetric the farther the flank and zero crossings become. This type of method is supposed to capture elements of the waveform that would be distorted by band pass filtering since it assumes sinusoidal basis functions, which are symmetric.

ence between adjacent extrema. The phase is assigned 0 at a maxima, π at a minima, $\frac{\pi}{2}$ at a negative flank point, and $\frac{3\pi}{2}$ at a positive flank point.

For burst detection amplitude consistency, period consistency, and monotonicity are calculated. Amplitude consistency is the ratio between adjacent amplitude calculations. Period consistency is the ratio between adjacent period calculations. Lastly, monotonicity is the ratio of the number of samples that are increasing between consecutive minimum and a maximum or the number of samples that are decreasing between consecutive maximum and a minimum. A threshold is then applied to determine if a burst is occurring, because bursts would represent excursions from baseline amplitude, period, or monotonicity. This gives a list of events by when they start and how long they occur.

Amplitude is taken from the previous amplitude calculations. Frequency is the inverse of the period calculations. Phase is the same as above. To get a continuous estimate for phase that is not step wise the estimate for phase is unwrapped and linearly interpolated between each successive estimate.

Algorithm 3 Peak and Trough Method

Signal is bandpass filtered for the specified band, generating $S_{bp}(t)$
Signal is lowpass filtered with corner frequency based on $lowband$ generating $S_{low}(t)$
Using the zero crossings from $S_bp(t)$, $S_{low}(t)$ is divided into cycles
The maximum or minimum within each cycle is determined based on the type of zero crossing
The flank points are determined as the time at which the value is exactly half of the adjacent extrema
The amplitude is half the absolute difference between adjacent extrema
The period is the time between each of the specified points (flanks and extrema)
The amplitude consistency, $AC(t)$ is calculated and the ratio between consecutive amplitude estimates
The period consistency, $PC(t)$, is calculated and the ratio between consecutive period estimates for a given cycle
The monotonicity, $MON(t)$, is calculated as the ratio of number of samples between a minimum to a maximum are increasing or from a maximum to a minum are decreasing for consecutive extrema
Each of $AC(t)$, $PC(t)$, and $MON(t)$ are compared to separate thresholds where a burst is detected if all thresholds are passed
Combine bursts separated by time gap smaller than $maxdrop$
Remove events shorter than $maxglitch$
for each event, i, of M bursts **do**
    The amplitude over time, $A(t)$ is determined as the amplitude as shown previously
    The phase is over time, $\theta(t)$ is the interpolation of the following assignments: 0 at a maxima, π at a minima, $\frac{\pi}{2}$ at a negative flank point, and $\frac{3\pi}{2}$ at a positive flank point.
    The frequency, $f(t)$, is the inverse of the period as shown previously
end for
return List of M events with A_i, f_i, θ_i for each event

III.2.4 Cosine Template

This method is an extension of "Hilbert Magnitude" method that uses analytic magnitude for detection but cosine template fitting to get event amplitude, frequency, and phase[1]. This is a implementation of published work[1] modified for extracted the parameter values over the whole course of the burst. This method will be referred to as the "Cosine Template" algorithm.

This method is exactly the same as the "Hilbert Magnitude" method up until all the events must be characterized. The detection method has generated a list of events by start time and duration.

Each event is split into several stride points at which points, amplitude, phase offset, and period will be estimated. Windows determine how much of the signal to use for fitting. Then each section is fit using Equation III.5. This gives estimates at the stride points. At each point amplitude is estimated as the value of A. The frequency is estimated as $\frac{1}{T}$. And phase is estimated over the whole window as $\phi_0 + 2\pi\frac{t}{T}$. For amplitude and frequency, values are interpolated over the stride points. For phase the values are averaged since there are overlapping sections.

$$A\sin(\phi_0 + 2\pi\frac{t}{T}) \tag{III.5}$$

Algorithm 4 Cosine Template Method

Signal is bandpass filtered for the specified band, generating $S_{bp}(t)$
The analytic signal, $s_a(t)$ is obtained using the Hilbert transform
$A(t) = abs(S_a(t))$
$\theta(t) = angle(S_a(t))$
$f(t) = \frac{d\theta(t)}{dt}$
$A_{slow}(t)$ comes from $A(t)$ lowpass filtered with time constant τ_{DC}
$P_{slow}(t) = A_{slow}(t)^2$
$P_{fast}(t) = A(t)^2$
The ratio of $P_{fast}(t)$ to $P_{slow}(t)$ with a putative burst detected if the ratio exceeds the threshold dB_{peak}
Extend detected bursts by comparing to another threshold dB_{end} only extended previous bursts to this lower threshold
Combine bursts separated by time gap smaller than $maxdrop$
Remove events shorter than $maxglitch$
This gives a list of M burst events
for each event, i, of M bursts **do**
    Each event is split into points at which parameters will be estimated, J
    for each stride point, j in J **do**
        A section of the event for curve fitting is determined by the forward and backward windows $w_{forward}$ and $w_{backward}$
        The section of the original signal, $S(t)$ is fit using $A\sin(\phi_0 + 2\pi\frac{t}{T})$
        Values for A, ϕ_0, and T are obtained
    end for
    The amplitude $A(t)$ is constructed from the interpolation of each estimate of A of the event
    The frequency is constructed from the interpolation of each estimate of $\frac{1}{T}$ of the event
    The phase is constructed from the interpolation of each estimate of $\phi_0 + 2\pi\frac{t}{T}$ of the event and averaged over overlapping sections
end for
return List of M events with A_i, f_i, θ_i for each event

III.2.5 Wavelet

Gabor-Morlet wavelet algorithms perform time-frequency decomposition to get magnitude and phase as a function of both time and frequency, expressing the signal as a sum of Gaussian-modulated sine waves [55]. Implementations differ depending on whether the Gaussian envelope has constant time duration over all frequencies or lasts for a constant number of cycles, constant-Q transform [70]. This work uses the continuous wavelet transform. This method will be referred to as the "Wavelet" algorithm.

This method starts by using either the constant Q transform `cqt` or the continuous wavelet transform `cwt` from MATLAB's Wavlet Library. The continuous wavelet transform with Morse wavelets was used for this method. Then the power is calculated from that, shown in Figure III.3. The power is then thresholded, shown in Figure III.4. The time values and frequency values are obtained for the portion above threshold. Opening and closing are both morphology functions found in image processing that are used in this method. Opening removes any features smaller than the given input structure or shape. Closing connects any features that are smaller than the given input structure or shape. Next opening is performed using an ellipse structuring element because the values in the x direction represent time and y direction represent frequency so a symmetric structuring element is not necessarily the best choice. Next border removal is performed, which removes any clusters that are connected to the borders. This removes extraneous above threshold detections and also removes any detections that may be at the border of a band if the analysis is done in a band. Lastly, closing is performed using an ellipse structuring element. Figure III.4 shows the resulting detection after each step. The groups are segmented into different blobs. These values are averaged for the center frequency and center time of the burst. The maximum and minimum frequencies are used as the burst edges, shown in Figure III.3. To get a more accurate frequency measurement, frequency is calculated within each time period of the detected burst by either taking the median, shown in Figure III.3 as the magenta line, or the power weighted average of the frequencies in the time slice. The detected power is calculated for each time slice as well as the summation of the magnitude values, where power values are converted to magnitude, summed, then converted back to power. The blue line in Figure III.3 represents the true burst. The frequency and burst power are also displayed in Figure III.5.

Algorithm 5 Wavelet Method

Spectrogram, S of signal is obtained using wavelet transform
$S_{mag} = abs(S)$
$S_{ang} = angle(S)$
Median Spectrogram, S_{med} obtained as median value of S per frequency
$S_{rel} = S_{mag} - S_{med}$
Convert S_{rel} into binary mask, S_{mask}, by threshold
Apply opening on S_{mask} with structuring element s_{open}
Remove foreground points connected to border of S_{mask}
Apply closing on S_{mask} with structuring element s_{close}
Isolate connected sections of S_{mask}
for each of M isolated sections of S_{mask} **do**
    S_{maski} is the mask in S_{mask} represented by the i section
    Apply mask S_{maski} to S to get S_i
    for Each sample of S_i **do**
        A_{in}, f_{in}, and θ_{in} is amplitude, frequency and phase of the n sample of section i
        $A_{in} = \max S_{mag_{in}}$
        $f_{in} = (S_{freq_{in}} * S_{mag_{in}}) / \sum S_{mag_{in}}$
        $\theta_{in} = S_{ang_{in}}$ at which $\max S_{mag_{in}}$ occurs
    end for
    A_i is the set of A_{in}
    f_i is the set of f_{in}
    θ_i is the set of θ_{in}
end for
return List of M events with A_i, f_i, θ_i for each event

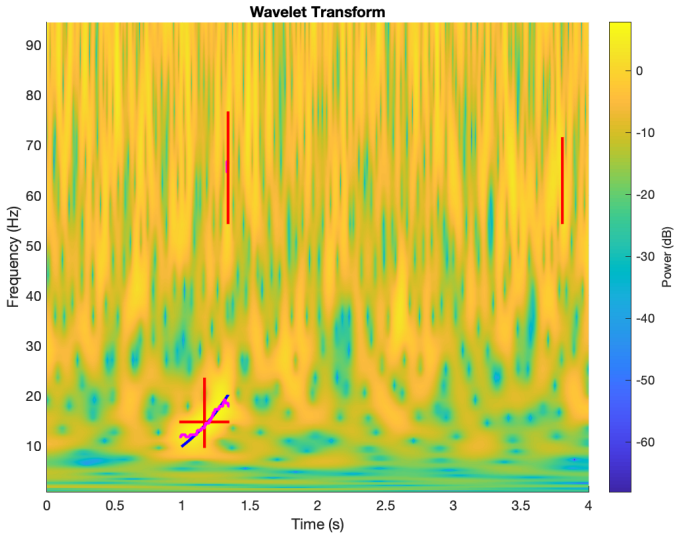

Figure III.3: Example detection with the wavelet method showing the spectrogram with detected burst indicated. The method outlines the start and end of each detected burst and the maximum and minimum frequency in each burst in red. The true burst is denoted in blue and the detected burst is denoted in magenta. The detected burst follows the true burst where it has highest amplitude in the middle of the burst and trails off at the ends.

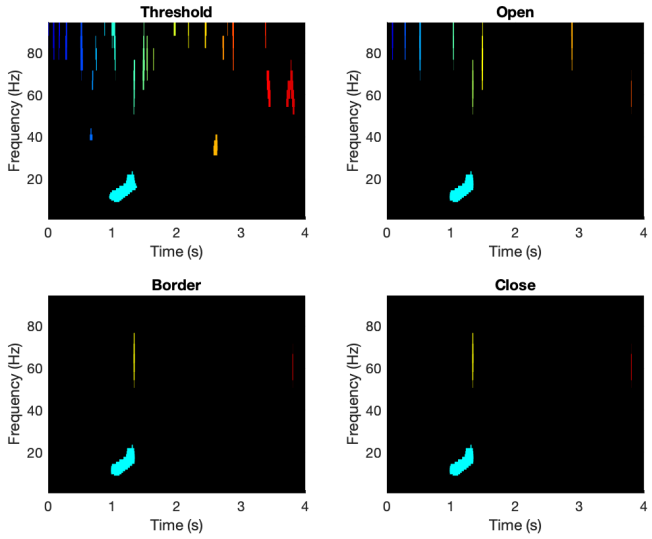

Figure III.4: Example detection with the wavelet method showing how the detection proceeds after getting the spectrogram from the wavelet transform. The top left plot shows the values in the spectrogram that are above the threshold and distinguishes them as separate groupings if they are unconnected. The top right plot shows the results after applying opening with an elliptical support. This process removes above threshold sections smaller than the structuring element, but preserving the sections that are larger than the structuring element. The bottom left plot shows the results after removing any sections that are connected to the border. This removes any events that may be distortions from either the beginning or end of the signal or the band pass. The bottom left plot shows the results after applying closing with an elliptical support. This process fills in sections that are closer together than the size of the structuring element.

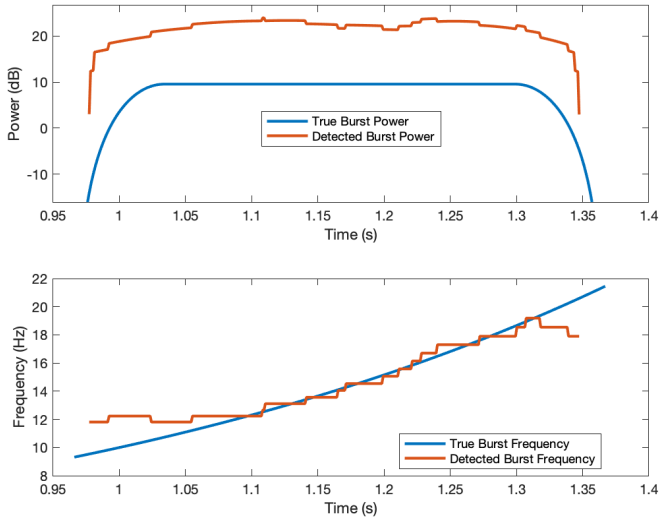

Figure III.5: Example detection with the wavelet method showing true burst parameters with detected burst parameters. The top plot shows that the amplitude detection of the detected burst is higher than the ground truth burst. This could be from the added power of the noise which increases the detected burst power. The bottom plot shows the frequency of the ground truth and detected bursts. The frequency tracks well in the center of the burst and trails off near the edges.

CHAPTER IV

Metrics for Performance Evaluation

IV.1 Introduction

The goal of this chapter is to evaluate the performance of detection and characterization methods in a meaningful way. This allows for comparison between different methods and will be used to indicate the effect of tuning. The metrics used are listed and described in the following sections. They are examined in detail for their use in tuning and ability to distinguish true and false detection. A good metric for tuning should provide information distinct from the others, demonstrate a search space with a well-behaved optimum, and represent a desirable feature.

To measure the detection performance the methods are compared to a known ground truth in the synthetic data set. Bursts that are identified by the selected method are compared to a list of events that are known to be in the signal. Matches between the ground truth list and the output of the method are True Positive (TP). A False Negative (FN) occurs when a ground truth event is not matched to a detected event from the method. The noise in the signal may cause events to be flagged that do not match any ground truth in the dataset leading to a False Positive (FP). There are no True Negatives (TN) because there are not single events that are not bursts in the ground truth sets. These counts make up the confusion matrix. Furthermore, additional detection metrics exist that are based on this count, including precision, recall, false positive rate, $\mathbf{F_1}$ Score,and $\mathbf{F}_\beta$ Score.

To measure the characterization performance the methods are compared to a baseline, which in this case is the band pass waveform and the acausal Hilbert derived magnitude, frequency, and phase unless otherwise noted[22]. None of the methods exactly output these values, but they do use them to construct their own characterizations. With this given baseline error can be calculated as the difference between the output of the selected method and the baseline. Value, shape, distribution of error values indicate various aspects of the error that may differentiate the performance of methods. The characterization metrics are a way to convert the vector of errors into a scalar value measuring some aspect of the overall error. Metrics dependent on the phase error come from circular statistics, namely mean direction, circular variance, and phase-locking value. Metrics for the remaining errors of signal, magnitude, and frequency are computed as the root mean square of the error and the relative error.

To obtain a metric that better encompasses the phase error in both the overall phase offset and the consistency of the error, the combined angle metric is proposed. This metric combines mean direction and circular variance.

IV.2 Confusion Matrix

Metrics based on the confusion matrix are used to measure the efficacy of the detection. This gives counts for True Positives, False Positives, and False Negatives. These are aggregated to obtain false positive rate, precision (Equation IV.5) and recall (Equation IV.9) can be calculated which allows for a relative comparison across methods. Combining these even further results in the F_1 (Equation IV.13)and F_β (Equation IV.18) scores, which encapsulate correctly identifying events and identifying as many true events as possible. These detection metrics will be called the confusion matrix family

A detected event is labeled as a True Positive if there is a matching ground truth event (as seen in Fig. IV.1a) and a False Positive otherwise (as seen in Fig. IV.1b); a ground truth event is labeled as a False Negative if there is no matching detected event (as seen in Fig. IV.1c).

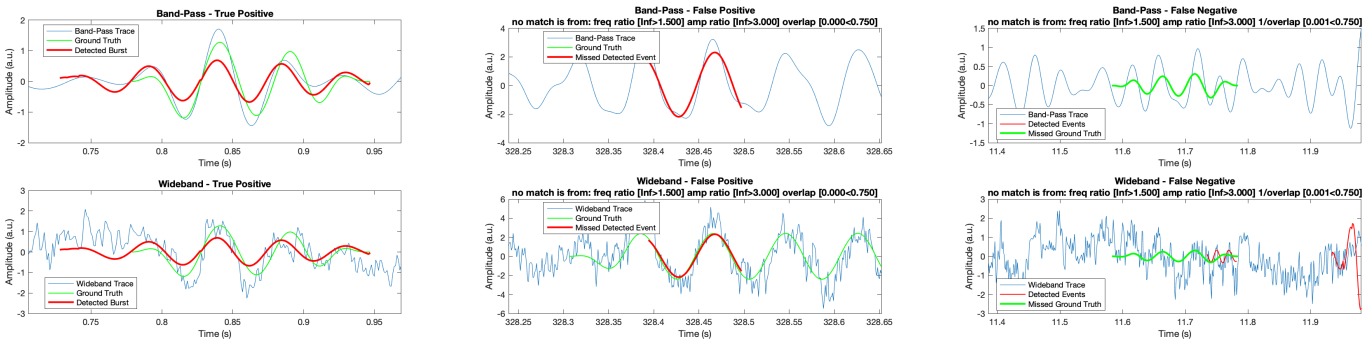

(a) True Positive Example

(b) False Positive Example

(c) False Negative Example

Figure IV.1: Examples of True Positive, False Positive, and False Negative Burst Detections.For each column, the top panel shows the detected and ground truth events relative to the band pass signal and the bottom panel shows the detected and ground truth events relative to the wide band signal. In (a) since there is a detected event that matches with the ground truth event this is labeled as a True Positive. In (c) the top panel shows that the ground truth event has no detected event in its band that matched. The bottom panel shows that there is not wide band detected event that matches the ground truth event. Since there is no matching detected event for the ground truth it is labeled as a False Negative. In (b) the top panel shows that there is no ground truth event in the band pass. The bottom panel shows that there is a potential corresponding ground truth event, but it does not match up with the detected event. The detected event is too short and is therefore classified as a False Positive.

IV.3 Detection Metrics

Detection metrics come from the confusion matrix information. They combine True Positive, False Positive, and False Negative to evaluate the detection.

To determine the error in the detection metrics the confusion matrix counts are modeled as Poisson random variables with parameter equal to the given count. This is because this closely models count data. Thus, the generation of an event in the synthetic data is a Poisson process, which is assumed to be similar in recorded data. From this the

True Positive, False Positive, and False Negative are also Poisson distributed since they represent a proportion of a Poisson process [71].

$$\delta_{FP} = \sqrt{FP} \tag{IV.1}$$

$$\delta_{TP} = \sqrt{TP} \tag{IV.2}$$

$$\delta_{FN} = \sqrt{FN} \tag{IV.3}$$

Because these counts are distributed as Poisson random variables their error, δ is represented accordingly in Equations IV.1, IV.2, and IV.3.

For combinations of these metrics error propagation must be done. In order to propagate the error Equation IV.4 shows error propagation for a function with three inputs.

$$\begin{aligned} f(x,y,z) \\ \delta_{fx} &= \left|\frac{\partial f}{\partial x}\right| \delta_x \\ \delta_{fy} &= \left|\frac{\partial f}{\partial y}\right| \delta_y \\ \delta_{fz} &= \left|\frac{\partial f}{\partial z}\right| \delta_z \\ \delta_f &= \sqrt{\delta_{fx}^2 + \delta_{fy}^2 + \delta_{fz}^2} \end{aligned} \tag{IV.4}$$

The partial derivative of the function is taken with respect to each input to determine the contribution of each to the error. Then the square root of the sum of errors squared gives the combined error for the function. This error propagation will be used with the detection metrics to determine the error of each.

IV.3.1 Precision

Also called Positive Predictive Value (PPV), represents the proportion of true positives within the number of detected events. A value of 1 indicates that all detected events are true events. A value of 0.5 indicates that half of the detected events are true events. A value of 0 indicates that none of the detected events are true events. The precision or PPV is defined,

$$PPV = \frac{TP}{TP + FP} \tag{IV.5}$$

Its standard deviation is shown in IV.8.

$$\delta_{PPV,TP} = \left|\frac{\partial Precision}{\partial TP}\right| \delta_{TP} = \left|\frac{FP}{(FP+TP)^2}\right| \sqrt{TP} \tag{IV.6}$$

$$\delta_{PPV,FP} = \left|\frac{\partial Precision}{\partial FP}\right| \delta_{FP} = \left|-\frac{TP}{(FP+TP)^2}\right| \sqrt{FP} \tag{IV.7}$$

$$\delta_{PPV} = \sqrt{\delta_{PPV,TP}^2 + \delta_{PPV,FP}^2} = \sqrt{\frac{FP*TP}{(FP+TP)^3}} \tag{IV.8}$$

IV.3.2 Recall

Also called True Positive Rate (TPR), hit rate, or sensitivity, recall represents the proportion of true events that are detected. A value of 1 indicates that all true events were detected. A value of 0.5 indicates that half of true events were detected. A value of 0 indicates that none of the true events were detected. The true positive rate or TPR is defined,

$$TPR = \frac{TP}{TP+FN} \tag{IV.9}$$

Its standard deviation is shown in IV.12.

$$\delta_{TPR,TP} = \left|\frac{\partial TPR}{\partial TP}\right| \delta_{TP} = \left|\frac{FN}{(FN+TP)^2}\right| \sqrt{TP} \tag{IV.10}$$

$$\delta_{TPR,FN} = \left|\frac{\partial TPR}{\partial FN}\right| \delta_{FN} = \left|-\frac{TP}{(FN+TP)^2}\right| \sqrt{FN} \tag{IV.11}$$

$$\delta_{TPR} = \sqrt{\delta_{TPR,TP}^2 + \delta_{TPR,FN}^2} = \sqrt{\frac{FN*TP}{(FN+TP)^3}} \tag{IV.12}$$

IV.3.3 $\mathbf{F_1}$ Score

The $\mathbf{F_1}$ Score compares TP, FP, FN counts. A value of 1 indicates that all the true events were detected and all the detected events were true events. A value of 0.5 indicates that there are just as many false detections and missed events (FP,FN) as there are correctly detected events (TP). A value of 0 indicates that none of the true events were detected or none of the detected events were true events. This metric combines evaluating detected events and the true events. The $\mathbf{F_1}$ Score is defined,

$$\mathbf{F_1} = 2\frac{PPV \cdot TPR}{PPV + TPR} \tag{IV.13}$$

Its standard deviation is shown in IV.17.

$$\begin{aligned}
\delta_{F_1,FP} &= \left|\frac{\partial F_1}{\partial FP}\right|\delta_{FP} = \left|\frac{\partial F_1}{\partial PPV}\frac{\partial PPV}{\partial FP} + \frac{\partial F_1}{\partial TPR}\frac{\partial TPR}{\partial FP}\right|\delta_{FP} \\
&= \left|\left(\frac{2*TPR^2}{(PPV+TPR)^2}\right)\left(-\frac{TP}{(FP+TP)^2}\right) + \left(\frac{2*PPV^2}{(PPV+TPR)^2}\right)(0)\right|\sqrt{FP} \\
&= \left|\left(\frac{2\,(FP+TP)^2}{(FP+FN+2TP)^2}\right)\left(-\frac{TP}{(FP+TP)^2}\right)\right|\sqrt{FP} \\
&= \frac{2*TP\sqrt{FP}}{(FP+FN+2TP)^2} \qquad \text{(IV.14)} \\
\delta_{F_1,TP} &= \left|\frac{\partial F_1}{\partial TP}\right|\delta_{TP} = \left|\frac{\partial F_1}{\partial PPV}\frac{\partial PPV}{\partial TP} + \frac{\partial F_1}{\partial TPR}\frac{\partial TPR}{\partial TP}\right|\delta_{TP} \\
&= \left|\left(\frac{2*TPR^2}{(PPV+TPR)^2}\right)\left(\frac{FP}{(FP+TP)^2}\right) + \left(\frac{2*PPV^2}{(PPV+TPR)^2}\right)\left(\frac{FN}{(TP+FN)^2}\right)\right|\sqrt{TP} \\
&= \left|\left(\frac{2\,(FP+TP)^2}{(FP+FN+2TP)^2}\right)\left(\frac{FP}{(FP+TP)^2}\right) + \left(\frac{2\,(FN+TP)^2}{(FP+FN+2TP)^2}\right)\left(\frac{FN}{(TP+FN)^2}\right)\right|\sqrt{TP} \\
&= \frac{2\,(FN+FP)\,\sqrt{TP}}{(FP+FN+2TP)^2} \qquad \text{(IV.15)} \\
\delta_{F_1,FN} &= \left|\frac{\partial F_1}{\partial FN}\right|\delta_{FN} = \left|\frac{\partial F_1}{\partial PPV}\frac{\partial PPV}{\partial FN} + \frac{\partial F_1}{\partial TPR}\frac{\partial TPR}{\partial FN}\right|\delta_{FN} \\
&= \left|\left(\frac{2*TPR^2}{(PPV+TPR)^2}\right)(0) + \left(\frac{2*PPV^2}{(PPV+TPR)^2}\right)\left(\frac{-TP}{(TP+FN)^2}\right)\right|\sqrt{FN} \\
&= \left|\left(\frac{2\,(FN+TP)^2}{(FP+FN+2TP)^2}\right)\left(\frac{-TP}{(TP+FN)^2}\right)\right|\sqrt{FN} \\
&= \frac{2*TP\sqrt{FN}}{(FP+FN+2TP)^2} \qquad \text{(IV.16)} \\
\delta_{F_1} &= \sqrt{\delta^2_{F_1,FP} + \delta^2_{F_1,TP} + \delta^2_{F_1,FN}} \\
&= \sqrt{\left(\frac{2*TP\sqrt{FP}}{(FP+FN+2TP)^2}\right)^2 + \left(\frac{2\,(FN+FP)\,\sqrt{TP}}{(FP+FN+2TP)^2}\right)^2 + \left(\frac{2*TP\sqrt{FN}}{(FP+FN+2TP)^2}\right)^2} \\
&= 2\sqrt{\frac{TP\,(FP+FN)\,(FP+TP+FN)}{(FP+2*TP+FN)^4}} \qquad \text{(IV.17)}
\end{aligned}$$

IV.3.4 $\mathbf{F}_\beta$ Score

The $\mathbf{F}_\beta$ Score is similar to the $\mathbf{F_1}$ score, but weights the evaluation of the detected events and the true events differently. The $\mathbf{F}_\beta$ Score is defined,

$$\mathbf{F}_\beta = (1+\beta^2)\frac{PPV \cdot TPR}{\beta^2 PPV + TPR} \qquad \text{(IV.18)}$$

Its standard deviation is shown in IV.22.

$$
\begin{aligned}
\delta_{F_\beta,FP} &= \left|\frac{\partial F_\beta}{\partial FP}\right|\delta_{FP} = \left|\frac{\partial F_\beta}{\partial PPV}\frac{\partial Precision}{\partial FP} + \frac{\partial F_\beta}{\partial TPR}\frac{\partial TPR}{\partial FP}\right|\delta_{FP} \\
&= \left|\left(\frac{(1+\beta^2)TPR^2}{((\beta^2 PPV + TPR)^2}\right)\left(-\frac{TP}{(FP+TP)^2}\right) + \left(\frac{(1+\beta^2)\beta^2 PPV}{((\beta^2 PPV + TPR)^2}\right)(0)\right|\sqrt{FP} \\
&= \left|\left(\frac{(1+\beta^2)\,(FP+TP)^2}{(FP+\beta^2 FN+(1+\beta^2)TP)^2}\right)\left(-\frac{TP}{(FP+TP)^2}\right)\right|\sqrt{FP} \\
&= \frac{(1+\beta^2)TP\sqrt{FP}}{(FP+\beta^2 FN+(1+\beta^2)TP)^2}
\end{aligned} \tag{IV.19}
$$

$$
\begin{aligned}
\delta_{F_\beta,TP} &= \left|\frac{\partial F_\beta}{\partial TP}\right|\delta_{TP} = \left|\frac{\partial F_\beta}{\partial PPV}\frac{\partial PPV}{\partial TP} + \frac{\partial F_\beta}{\partial TPR}\frac{\partial TPR}{\partial TP}\right|\delta_{TP} \\
&= \left|\left(\frac{(1+\beta^2)TPR^2}{((\beta^2 PPV + TPR)^2}\right)\left(\frac{FP}{(FP+TP)^2}\right) + \left(\frac{(1+\beta^2)\beta^2 PPV}{((\beta^2 PPV + TPR)^2}\right)\left(\frac{FN}{(TP+FN)^2}\right)\right|\sqrt{TP} \\
&= \left|\left(\frac{(1+\beta^2)\,(FP+TP)^2}{(FP+\beta^2 FN+(1+\beta^2)TP)^2}\right)\left(\frac{FP}{(FP+TP)^2}\right) + \left(\frac{(1+\beta^2)\beta^2\,(FN+TP)^2}{(FP+\beta^2 FN+(1+\beta^2)TP)^2}\right)\left(\frac{FN}{(TP+FN)^2}\right)\right|\sqrt{TP} \\
&= \frac{(1+\beta^2)\left(\beta^2 FN+FP\right)\sqrt{TP}}{(FP+\beta^2 FN+(1+\beta^2)TP)^2}
\end{aligned} \tag{IV.20}
$$

$$
\begin{aligned}
\delta_{F_\beta,FN} &= \left|\frac{\partial F_\beta}{\partial FN}\right|\delta_{FN} = \left|\frac{\partial F_\beta}{\partial PPV}\frac{\partial PPV}{\partial FN} + \frac{\partial F_\beta}{\partial TPR}\frac{\partial TPR}{\partial FN}\right|\delta_{FN} \\
&= \left|\left(\frac{(1+\beta^2)TPR^2}{((\beta^2 PPV + TPR)^2}\right)(0) + \left(\frac{(1+\beta^2)\beta^2 PPV}{((\beta^2 PPV + TPR)^2}\right)\left(\frac{-TP}{(TP+FN)^2}\right)\right|\sqrt{FN} \\
&= \left|\left(\frac{(1+\beta^2)\beta^2\,(FN+TP)^2}{(FP+\beta^2 FN+(1+\beta^2)TP)^2}\right)\left(\frac{-TP}{(TP+FN)^2}\right)\right|\sqrt{FN} \\
&= \frac{(1+\beta^2)\beta^2 TP\sqrt{FN}}{(FP+\beta^2 FN+(1+\beta^2)TP)^2}
\end{aligned} \tag{IV.21}
$$

$$
\begin{aligned}
\delta_{F_\beta} &= \sqrt{\delta^2_{F_\beta,FP} + \delta^2_{F_\beta,TP} + \delta^2_{F_\beta,FN}} \\
&= \sqrt{\left(\frac{(1+\beta^2)TP\sqrt{FP}}{(FP+\beta^2 FN+(1+\beta^2)TP)^2}\right)^2 + \left(\frac{(1+\beta^2)\,(\beta^2 FN+FP)\sqrt{TP}}{(FP+\beta^2 FN+(1+\beta^2)TP)^2}\right)^2 + \left(\frac{(1+\beta^2)\beta^2 TP\sqrt{FN}}{(FP+\beta^2 FN+(1+\beta^2)TP)^2}\right)^2} \\
&= (1+\beta^2)\sqrt{\frac{TP\,(FP^2+FN\,(FN+TP)\,\beta^4+FP\,(TP+2*FN\beta^2))}{(FP+\beta^2 FN+(1+\beta^2)TP)^4}}
\end{aligned} \tag{IV.22}
$$

IV.4 Characterization Metrics

For characterization metrics, the residual is used. The residual is defined here as the difference between the i-th extracted feature of the j-th event and a reference for the i-th feature of the j-th event as shown in Equation IV.23.

$$
e_{i,j}(t) = Ext_{i,j}(t) - Ref_{i,j}(t) \tag{IV.23}
$$

This residual between the detected feature and a reference feature are computed for signal, magnitude, frequency, and phase. Equation IV.24 gives the residual for the m-th event for the signal, magnitude, frequency, and phase represented by the indices s, m, f, and ϕ respectively. Here $e_{s,m}(t)$ represents the signal residual of the given event, m, for the given method described by $\Psi_m(A_m(t), f_m(t), \phi_{0_m}, t_{0_m})$. $N_A(t)$. The subsequent residual terms represent the calculations for magnitude, frequency, and phase residual.

$$
\begin{aligned}
e_{s,m}(t) &= \Psi_m(A_m(t), f_m(t), \phi_{0_m}, t_{0_m}) - N_A(t) \\
e_{m,m}(t) &= A_m(t) - H_m(N_A(t)) \\
e_{f,m}(t) &= f_m(t) - H_f(N_A(t)) \\
e_{\phi,m}(t) &= \phi_m(t) - H_\phi(N_A(t))
\end{aligned} \tag{IV.24}
$$

In the following two sections the residuals for signal, magnitude, and frequency are treated with RMS calculations (per section IV.5), and phase, which is defined on a circular domain, is treated with circular statistics (per section IV.6)

IV.5 RMS Metrics

For signal, magnitude, and frequency residual values, one metric is to take the Root Mean Square (RMS), Equation IV.25, of the residuals from Equation IV.24. Here the index, i, represetns the type of residual used (signal, magnitude, or frequency) and the index, j, represents the event overall. For the j-th event it has N samples. Additionally for each a relative error can be calculated, whereby the value of the error is scaled by the value at that point from the reference method. This is so that errors at very large values and errors at very small values will have proportional effects on the error metric. For the relative signal value, the error associated with points near a zero crossing of the sinusoid are either removed or interpolated with both implemented.

$$
RMS_{i,j} = \sqrt{\frac{1}{N}\sum_{n=1}^{N} e_{i,j}(n)^2} \tag{IV.25}
$$

In this section ten RMS signal metrics are explored. They all are functions of the residuals and incorporate the RMS.

1. Signal Metric
2. Magnitude Metric
3. Frequency Metric
4. Relative Event Power to Bandpass Metric
5. Relative Event Power to Wideband Metric
6. Relative Signal Metric

7. Removed Relative Signal Metric
8. Interpolated Relative Signal Metric
9. Relative Magnitude Metric
10. Relative Frequency Metric

IV.5.1 Signal Metric

Signal Metric is the RMS of the instantaneous signal residual. This is calculated using Equation IV.25 with signal residuals. This is a measure of how much the estimated burst signal deviates from a reference signal. In Figure IV.2 two events with different Signal Metric values are shown. Figure IV.2a shows a low value for the Signal Metric which is indicative of an event that closely follows the reference signal. This incorporates a good estimate of amplitude, phase, and frequency to match the reference well. Figure IV.2b shows a high value for the Signal Metric, showing an event that does not correspond well with the reference signal indicating a poorer characterization of the event.

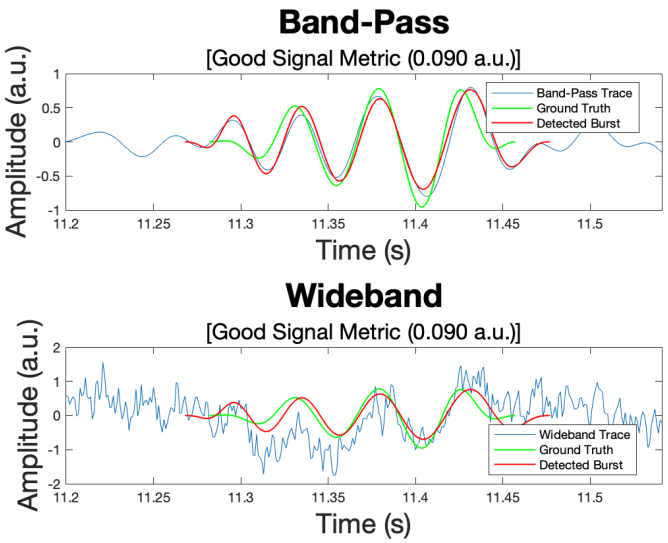

(a) Good Signal Metric

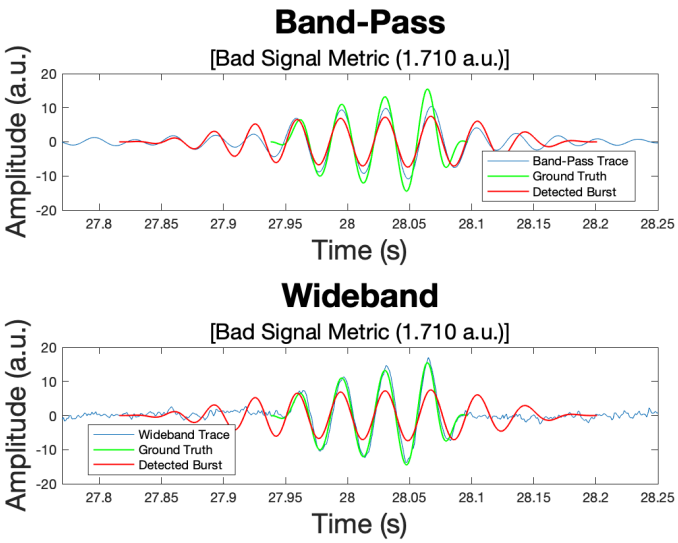

(b) Bad Signal Metric

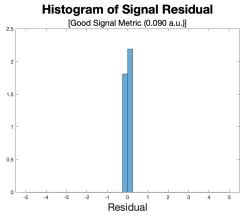

(c) Good Signal Metric Histogram of Residuals

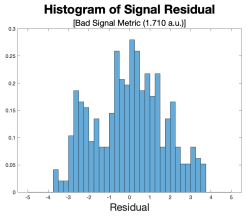

(d) Bad Signal Metric Histogram of Residuals

Figure IV.2: Here examples of good and bad event characterizations are shown as determined by Signal Metric. In (a) a low value of Signal Metric is shown indicating an event that closely follows the reference event. In (b) a high value of Signal Metric is shown indicating an event that does not closely follow the reference event.

Figure IV.3 shows how well this metric does at separating events by whether they represent a true or false detection. Figure IV.3a shows how the metric separates the true or false detections. The detections cluster around each other indicating that this metric cannot be used on its own to filter true detections from false detections without rejecting many true detections. Figure IV.3b shows how the metric separates the true and false detections along with how it separates based on the signal to noise ratio. While this figure shows that low SNR events are not well detected which is expected, it does not show that the Signal Metric does not provide a useful metric for separating the events by true and false detections without rejecting many true detections. A cutoff around 0.2 will reject almost all false detections and also reject many true detections.

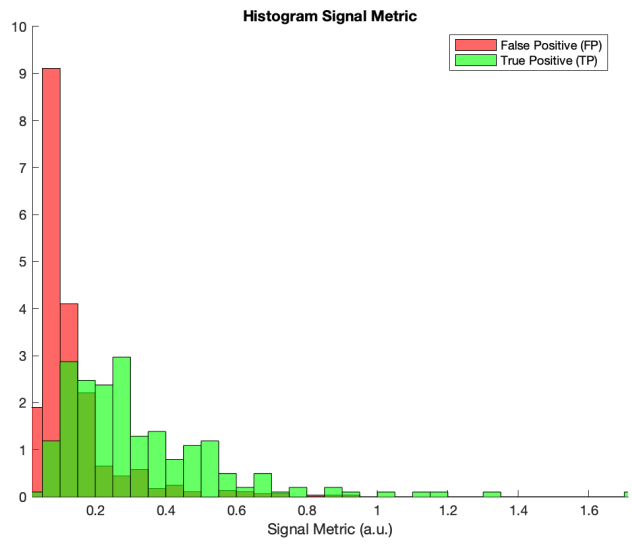

(a) Histogram Signal Metric

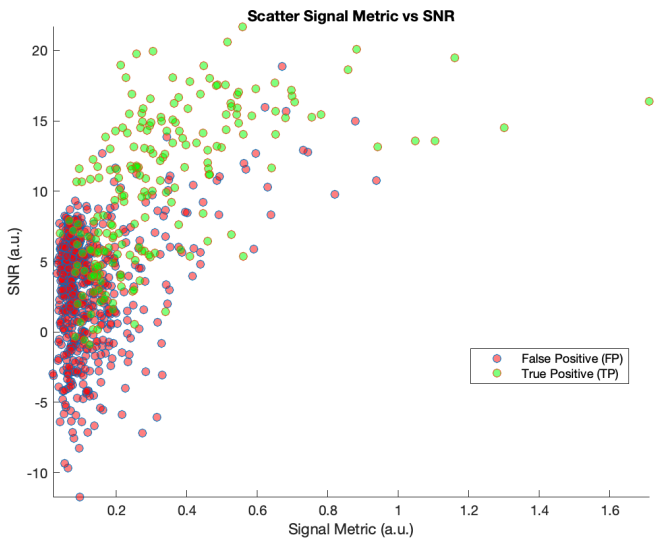

(b) Signal Metric Vs SNR

Figure IV.3: This figure shows the types of events, either true positive, false positive, or false negative, that the Magnitude Method generates. These events are shown in a histogram in (a) are binned by the Signal Metric. These events are also shown in a scatter plot in (b) which are separated by the Signal Metric and the signal to noise ratio. These plots show how the metric could be used for determining true and false detections.

IV.5.2 Magnitude Metric

Magnitude Metric is the RMS of the instantaneous magnitude residual. This is calculated using Equation IV.25 with magnitude residuals. The Magnitude Metric is a measure of how much the estimated burst magnitude deviates from a reference magnitude. In Figure IV.4 two events with different Magnitude Metric values are shown. Figure IV.4a shows a low value for the Magnitude Metric which is indicative of an event that closely follows the reference magnitude. Figure IV.4b shows a high value for the Magnitude Metric, showing an event that does not correspond well with the reference magnitude indicating a poorer characterization of the event magnitude.

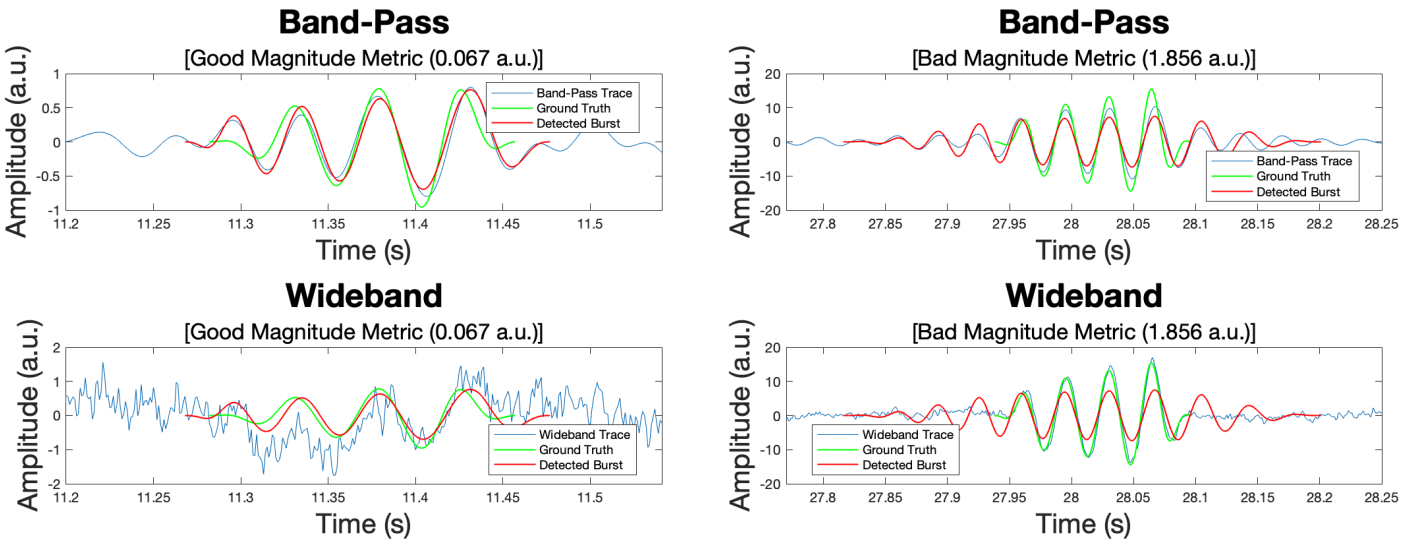

(a) Good Magnitude Metric

(b) Bad Magnitude Metric

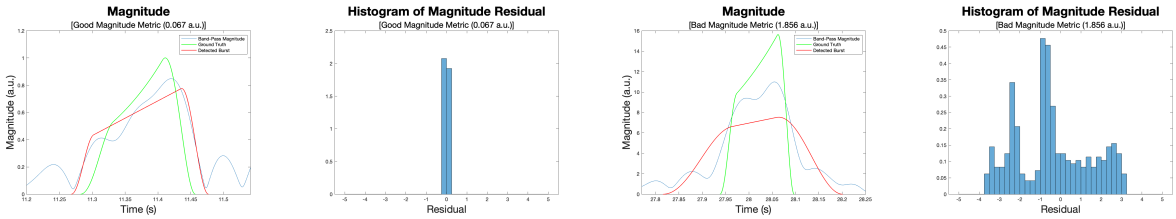

(c) Good Magnitude Metric Magnitude Plot

(d) Good Magnitude Metric Histogram of Residuals

(e) Bad Magnitude Metric Magnitude Plot

(f) Bad Magnitude Metric Histogram of Residuals

Figure IV.4: Here examples of good and bad event characterizations are shown as determined by Magnitude Metric. In (a) a low value of Magnitude Metric is shown indicating an event that closely follows the reference. In (b) a high value of Magnitude Metric is shown indicating an event that does not closely follow the reference. Both (c) and (e) show the extracted magnitude in red and how it compares to the reference in blue, with the addition of ground truth events in green. Both (d) and (f) show the residuals for the detected events.

Figure IV.5 shows how well this metric does at separating events by whether they represent a true or false detection. Figure IV.5a shows how the metric separates the true or false detections. The detections cluster around each other indicating that this metric cannot be used on its own to filter true detections from false detections without rejecting some true detections. Figure IV.5b shows how events separates along the metric and the signal to noise ratio. While this figure shows that low SNR events are typically false positives, it does not show that the Magnitude Metric is a useful metric for separating the events by true and false detections without rejecting some true detections. A cutoff around 0.15 will reject almost all false detections and also reject some true detections.

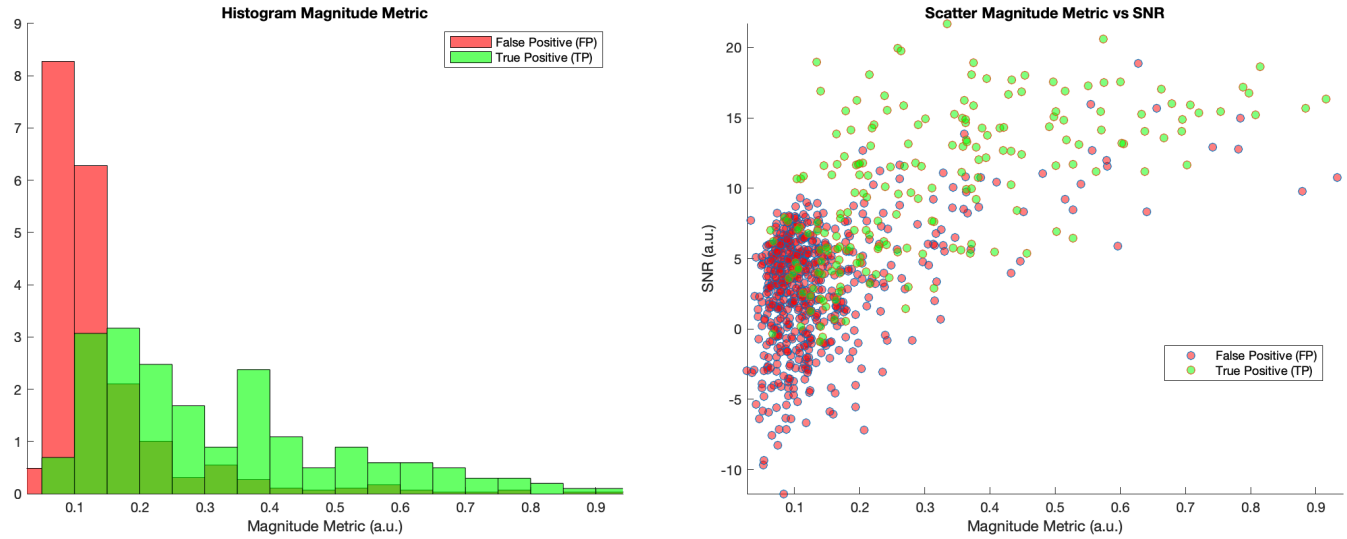

(a) Histogram Magnitude Metric

(b) Magnitude Metric Vs SNR

Figure IV.5: This figure shows the types of events, either true positive, false positive, or false negative, that the Magnitude Method generates. These events are shown in a histogram in (a) are binned by the Magnitude Metric. These events are also shown in a scatter plot in (b) which are separated by the Magnitude Metric and the signal to noise ratio. These plots show how the metric could be used for determining true and false detections.

IV.5.3 Frequency Metric

Frequency Metric is the RMS of the instantaneous frequency residual. This is calculated using Equation IV.25 with frequency residuals. The Frequency Metric is a measure of how much the estimated burst frequency deviates from a reference frequency. In Figure IV.6 two events with different Frequency Metric values are shown. Figure IV.6a shows a low value for the Frequency Metric which is indicative of an event that closely follows the reference frequency. Figure IV.6b shows a high value for the Frequency Metric, showing an event that does not correspond well with the reference frequency indicating a poorer characterization of the event frequency.

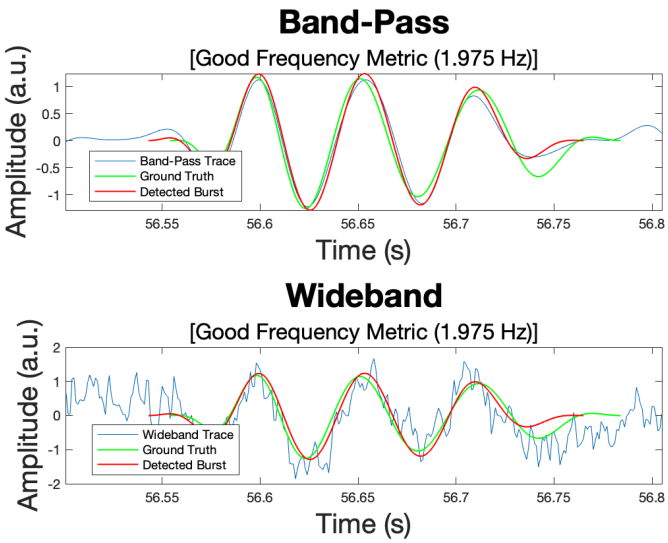

(a) Good Frequency Metric

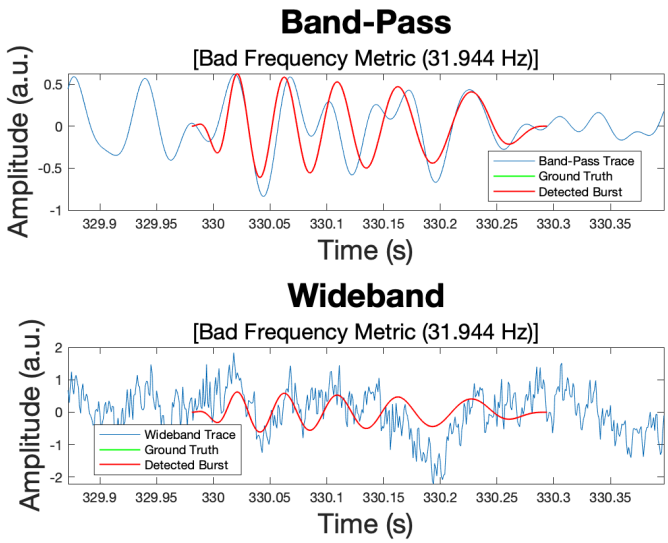

(b) Bad Frequency Metric

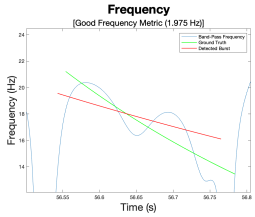

(c) Good Frequency Metric Frequency Plot

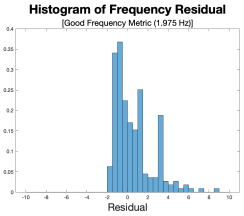

(d) Good Frequency Metric Histogram of Residuals

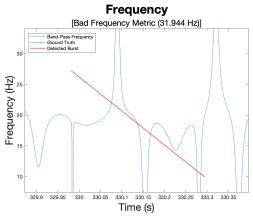

(e) Bad Frequency Metric Frequency Plot

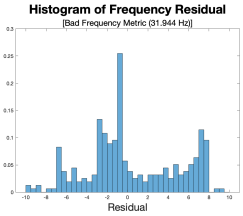

(f) Bad Frequency Metric Histogram of Residuals

Figure IV.6: Here examples of good and bad event characterizations are shown as determined by Frequency Metric. In (a) a low value of Frequency Metric is shown indicating an event that closely follows the reference. In (b) a high value of Frequency Metric is shown indicating an event that does not closely follow the reference. Both (c) and (e) show the extracted frequency in red and how it compares to the reference in blue, with the addition of ground truth events in green. Both (d) and (f) show the residuals for the detected events.

Figure IV.7 shows how well this metric does at separating events by whether they represent a true or false detection. Figure IV.7a shows how the metric separates the true or false detections. The detections cluster around each other indicating that this metric cannot be used on its own to filter true detections from false detections. Figure IV.7b shows how events separates along the metric and the signal to noise ratio. While this figure shows that low SNR events are typically false positives, it does not show that the Frequency Metric is a useful metric for separating the events by true and false detections.

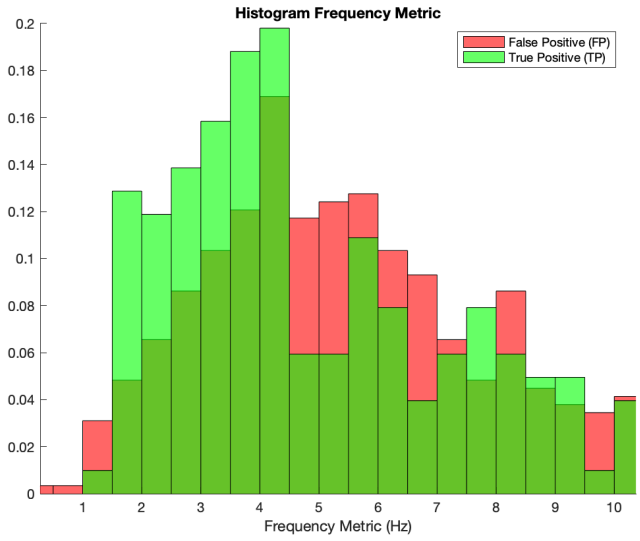

(a) Histogram Frequency Metric

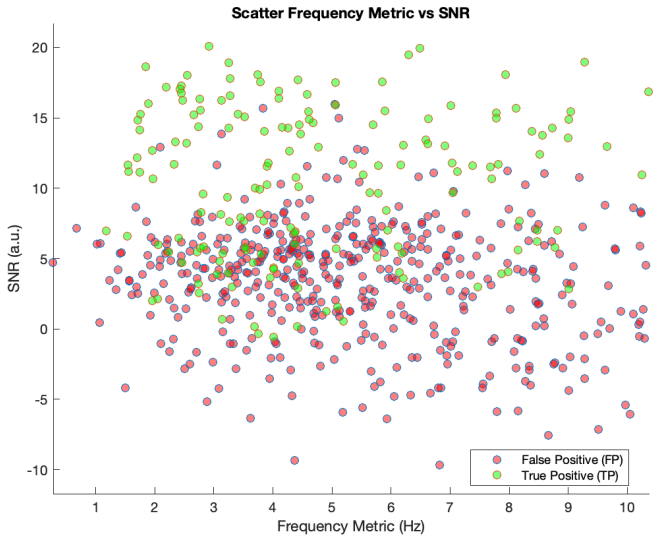

(b) Frequency Metric Vs SNR

Figure IV.7: This figure shows the types of events, either true positive, false positive, or false negative, that the Magnitude Method generates. These events are shown in a histogram in (a) are binned by the Frequency Metric. These events are also shown in a scatter plot in (b) which are separated by the Frequency Metric and the signal to noise ratio. These plots show how the metric could be used for determining true and false detections.

IV.5.4 Relative Event Power to Bandpass Metric

Relative Event Power to Bandpass Metric is the RMS of the signal residual relative to the RMS of bandpass signal. For an event, j, Equation IV.26 shows how relative power is calculated. Here The reference is the signal, either wideband or bandpass, with length N. Relative Event Power to Bandpass Metric is calculated using Equation IV.26 with signal residuals and the reference signal as the bandpass signal.

$$RelPow_j = \frac{RMS_{s,j}}{\sqrt{\frac{1}{N}\sum_{n=1}^{N} Ref_{s,j}(n)^2}} \tag{IV.26}$$

Relative Event Power to Bandpass Metric is a metric that uses the power of the reference signal and the power in the error signal. The power in the error signal is calculated as the RMS of the Signal Residual as above. The power in the reference signal is RMS of the bandpass of the signal. The ratio of the RMS of the Signal Residual to the RMS of the bandpass signal over the time period of the event is the Relative Event Power to Bandpass Metric.

The Relative Event Power to Bandpass Metric is a measure of how much the estimated burst signal matches the bandpass signal. This gives a measure of signal error that is relative to the power of the signal. Events with greater power might have larger RMS of the Signal Residual not because of larger error, but simply because of a larger magnitude. This metric would aim to scale that value to be able to compare events more fairly.

In Figure IV.8 two events with different Relative Event Power to Bandpass Metric values are shown. Figure IV.8a shows a low value for the Relative Event Power to Bandpass Metric which is indicative of an event that closely follows the reference. Figure IV.8b shows a high value for the Relative Event Power to Bandpass Metric, showing an event that does not correspond well with the reference indicating a poorer characterization of the event.

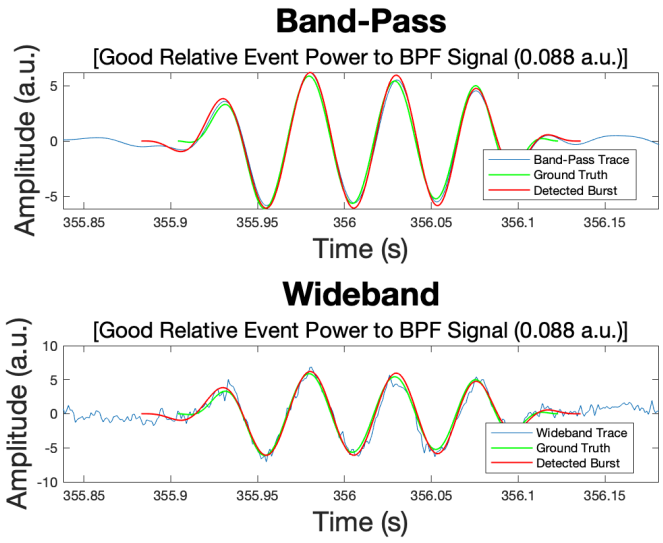

(a) Good Relative Event Power to Bandpass Metric

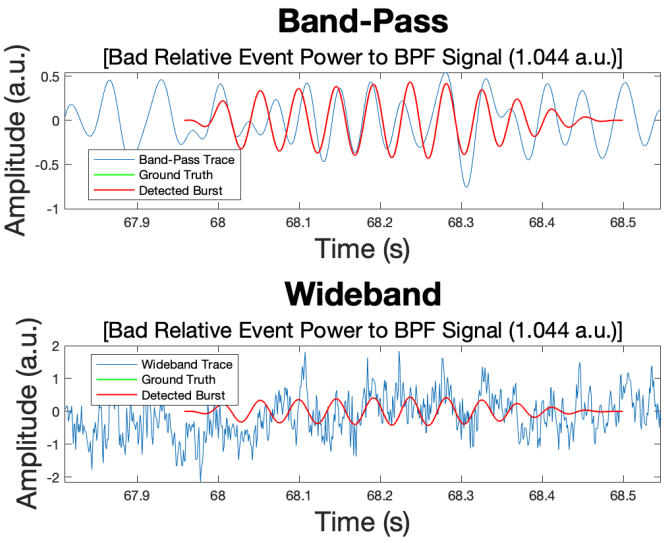

(b) Bad Relative Event Power to Bandpass Metric

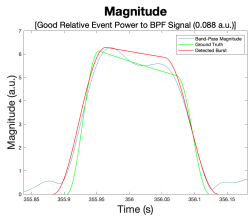

(c) Good Relative Event Power to Bandpass Metric Magnitude Plot

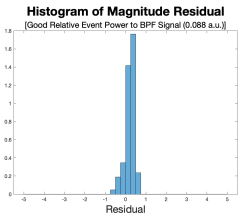

(d) Good Relative Event Power to Bandpass Metric Histogram of Residuals

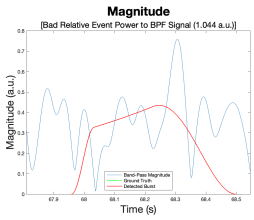

(e) Bad Relative Event Power to Bandpass Metric Magnitude Plot

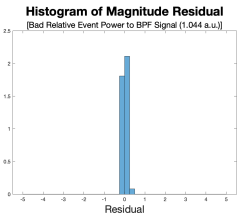

(f) Bad Relative Event Power to Bandpass Metric Histogram of Residuals

Figure IV.8: Here examples of good and bad event characterizations are shown as determined by Relative Event Power to Bandpass Metric. In (a) a low value of Relative Event Power to Bandpass Metric is shown indicating an event that closely follows the reference. In (b) a high value of Relative Event Power to Bandpass Metric is shown indicating an event that does not closely follow the reference. Both (c) and (e) show the extracted magnitude in red and how it compares to the reference in blue, with the addition of ground truth events in green. Both (d) and (f) show the residuals for the detected events.

Figure IV.9 shows how well this metric does at separating events by whether they represent a true or false detection. Figure IV.9a shows how the metric separates the true or false detections. The detections cluster around each other indicating that this metric cannot be used on its own to filter true detections from false detections. Figure

IV.9b shows how events separates along the metric and the signal to noise ratio. While this figure shows that low SNR events are typically false positives, it does not show that the Relative Event Power to Bandpass Metric is a useful metric for separating the events by true and false detections. The metric performs better as a filter with a cutoff of 0.25 for events with SNR above 10.

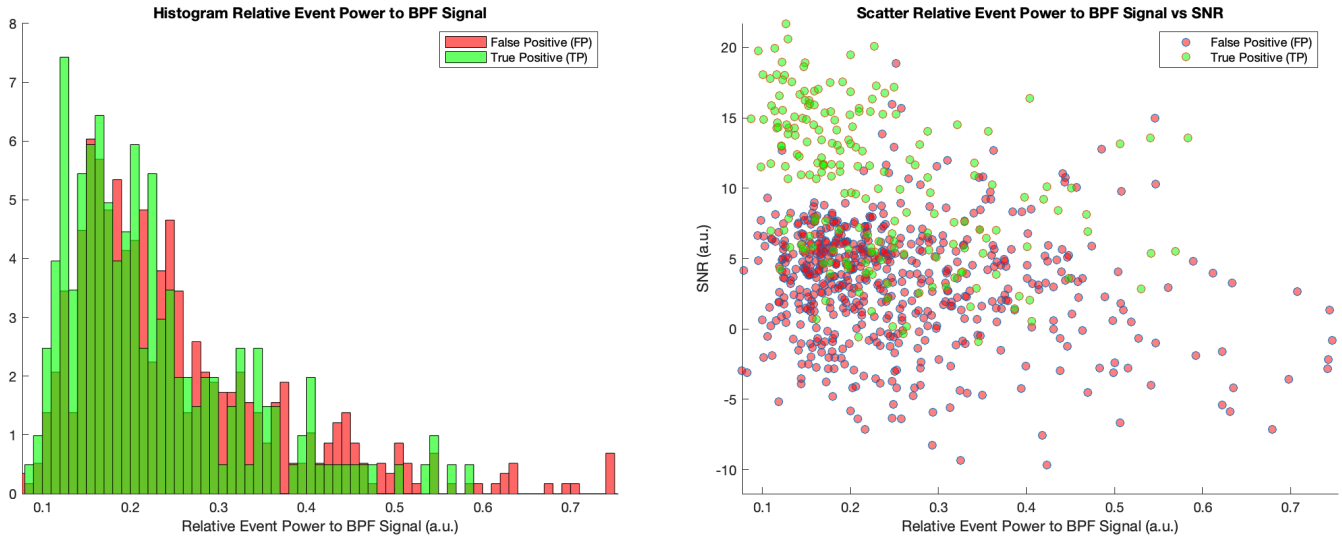

(a) Histogram Relative Event Power to Bandpass Metric

(b) Relative Event Power to Bandpass Metric Vs SNR

Figure IV.9: This figure shows the types of events, either true positive, false positive, or false negative, that the Magnitude Method generates. These events are shown in a histogram in (a) are binned by the Relative Event Power to Bandpass Metric. These events are also shown in a scatter plot in (b) which are separated by the Relative Event Power to Bandpass Metric and the signal to noise ratio. These plots show how the metric could be used for determining true and false detections.

IV.5.5 Relative Event Power to Wideband Metric

Relative Event Power to Wideband Metric is the RMS of the signal residual relative to the RMS of wideband signal. This is calculated using Equation IV.26 with signal residuals and the reference signal as the wideband signal

The Relative Event Power to Wideband Metric is a metric that uses the power of the reference signal and the power in the error signal. The power in the error signal is calculated as the RMS of the Signal Residual as above. The power in the reference signal is RMS of the wide band of the signal. The ratio of the RMS of the Signal Residual to the RMS of the wide band signal over the time period of the event is the Relative Event Power to Wideband Metric.

The Relative Event Power to Wideband Metric is a measure of how much the estimated burst signal matches the wide band signal. This gives a measure of signal error that is relative to the power of the signal. Events with greater power might have larger RMS of the Signal Residual not because of larger error, but simply because of a larger magnitude. This metric would aim to scale that value to be able to compare events more fairly.

In Figure IV.10 two events with different Relative Event Power to Wideband Metric values are shown. Figure IV.10a shows a low value for theRelative Event Power to Wideband Metric which is indicative of an event that closely follows the reference. Figure IV.10b shows a high value for the Relative Event Power to Wideband Metric, showing an event that does not correspond well with the reference indicating a poorer characterization of the event.

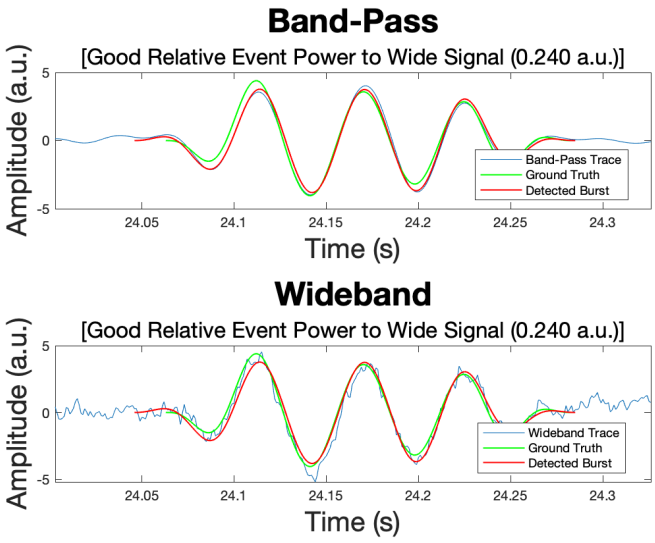

(a) Good Relative Event Power to Wideband Metric

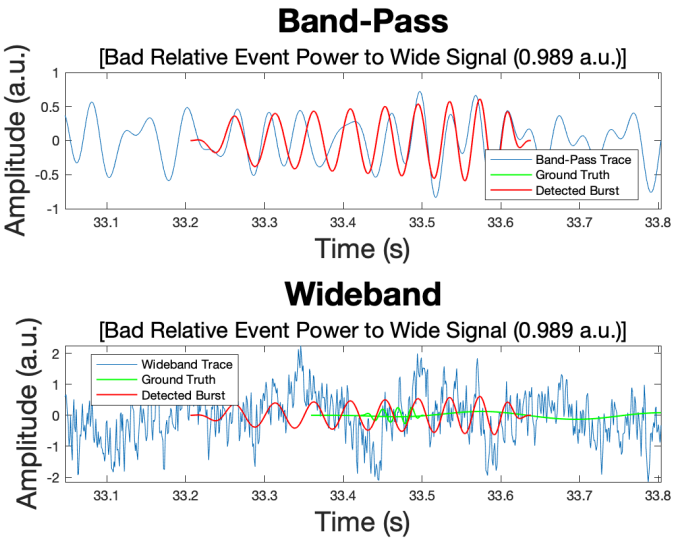

(b) Bad Relative Event Power to Wideband Metric

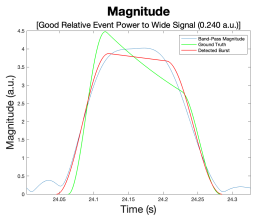

(c) Good Relative Event Power to Wideband Metric Magnitude Plot

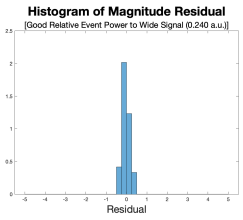

(d) Good Relative Event Power to Wideband Metric Histogram of Residuals

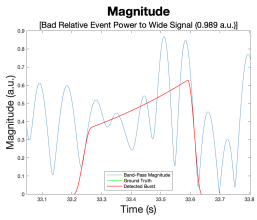

(e) Bad Relative Event Power to Wideband Metric Magnitude Plot

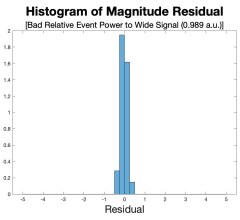

(f) Bad Relative Event Power to Wideband Metric Histogram of Residuals

Figure IV.10: Here examples of good and bad event characterizations are shown as determined by Relative Event Power to Wideband Metric. In (a) a low value of Relative Event Power to Wideband Metric is shown indicating an event that closely follows the reference. In (b) a high value of Relative Event Power to Wideband Metric is shown indicating an event that does not closely follow the reference. Both (c) and (e) show the extracted magnitude in red and how it compares to the reference in blue, with the addition of ground truth events in green. Both (d) and (f) show the residuals for the detected events.

Figure IV.11 shows how well this metric does at separating events by whether they represent a true or false detection. Figure IV.11a shows how the metric separates the true or false detections. The detections cluster around each other but not as much as previous metrics, indicating that this metric could provide useful information on its

own to filter true detections from false detections. Figure IV.11b shows how events separates along the metric and the signal to noise ratio. This figure shows that low SNR events are typically false positives, and also that the Relative Event Power to Wideband Metric is a possibly useful metric for separating the events by true and false detections. A cutoff around 0.7 will reject many false detections and also reject some true detections.

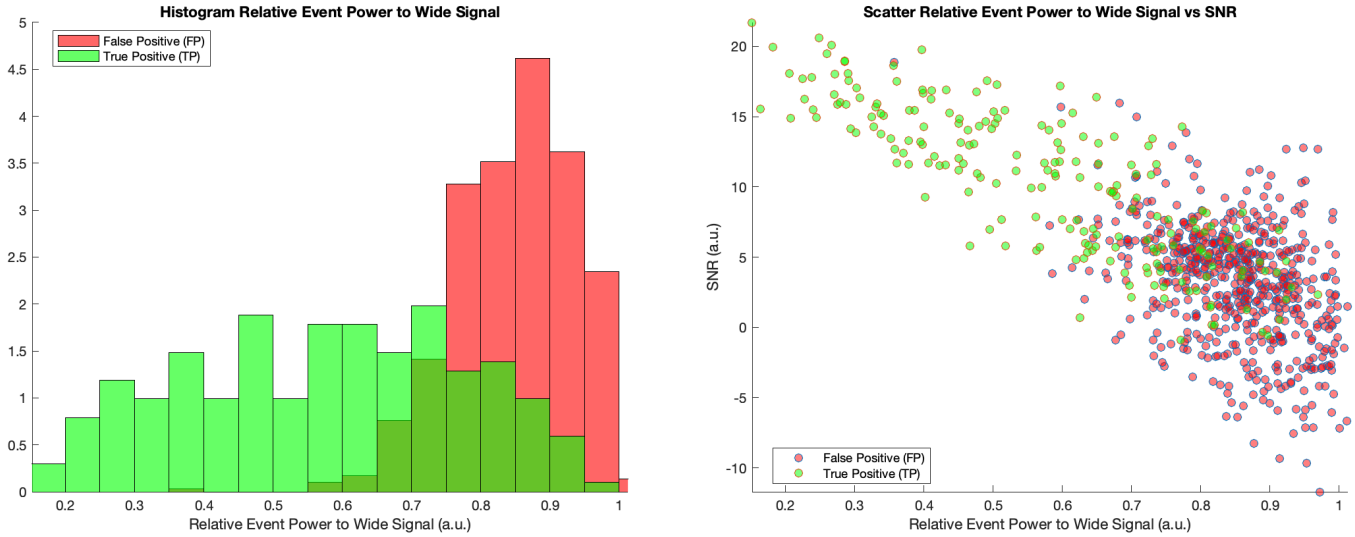

(a) Histogram Relative Event Power to Wideband Metric

(b) Relative Event Power to Wideband Metric Vs SNR

Figure IV.11: This figure shows the types of events, either true positive, false positive, or false negative, that the Magnitude Method generates. These events are shown in a histogram in (a) are binned by the Relative Event Power to Wideband Metric. These events are also shown in a scatter plot in (b) which are separated by the Relative Event Power to Wideband Metric and the signal to noise ratio. These plots show how the metric could be used for determining true and false detections.

IV.5.6 Relative Signal Metric

Relative Signal Metric is the RMS of the instantaneous relative signal residual. Equation IV.27 shows how the relative residual for the i-th feature for the j-th event is calculated. The Relative Signal Metric is calculated using Equation IV.25 with relative signal residuals.

$$\hat{e}_{i,j}(t) = \frac{Ext_{i,j}(t) - Ref_{i,j}(t)}{Ref_{i,j}(t)} \tag{IV.27}$$

Relative Signal Metric is a measure of how much the estimated burst signal deviates from a reference signal. In Figure IV.12 two events with different Relative Signal Metric values are shown. Figure IV.12a shows a low value for the Relative Signal Metric which is indicative of an event that closely follows the reference signal. This incorporates a good estimate of amplitude, phase, and frequency to match the reference well. Figure IV.12b shows a high value

for the Relative Signal Metric, showing an event that does not correspond well with the reference signal indicating a poorer characterization of the event.

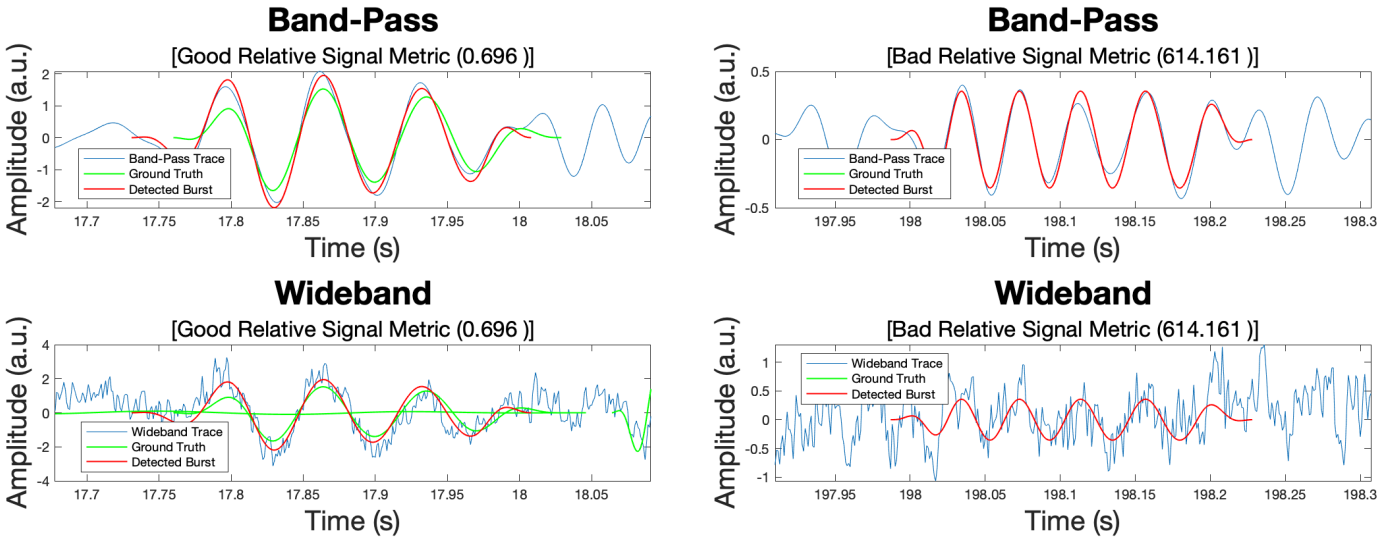

(a) Good Relative Signal Metric

(b) Bad Relative Signal Metric

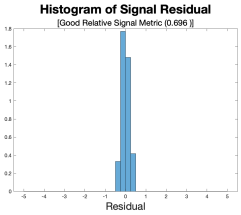

(c) Good Relative Signal Metric Histogram of Residuals

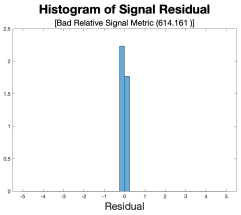

(d) Bad Relative Signal Metric Histogram of Residuals

Figure IV.12: Here examples of good and bad event characterizations are shown as determined by Relative Signal Metric. In (a) a low value of Relative Signal Metric is shown indicating an event that closely follows the reference. In (b) a high value of Relative Signal Metric is shown indicating an event that does not closely follow the reference. Both (c) and (d) show the residuals for the detected events.

Figure IV.13 shows how well this metric does at separating events by whether they represent a true or false detection. Figure IV.13a shows how the metric separates the true or false detections. The detections cluster around each other indicating that this metric cannot be used on its own to filter true detections from false detections. Figure IV.13b shows how events separates along the metric and the signal to noise ratio. While this figure shows that low SNR events are typically false positives, it does not show that the Relative Signal Metric is a useful metric for separating the events by true and false detections.

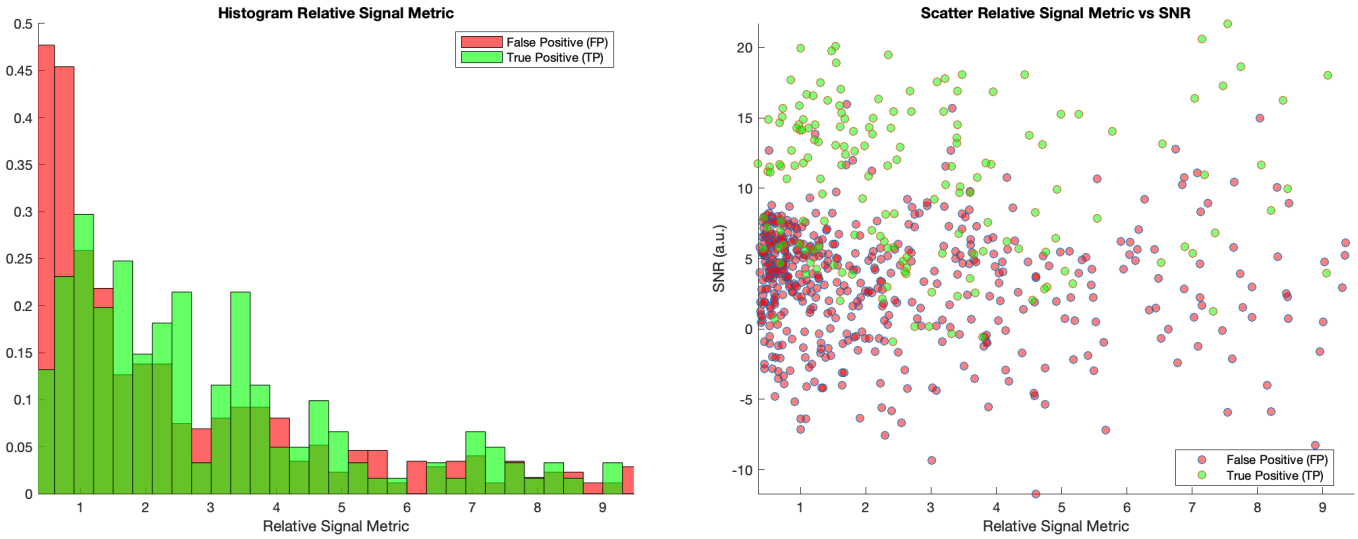

(a) Histogram Relative Signal Metric

(b) Relative Signal Metric Vs SNR

Figure IV.13: This figure shows the types of events, either true positive, false positive, or false negative, that the Magnitude Method generates. These events are shown in a histogram in (a) are binned by the Relative Signal Metric. These events are also shown in a scatter plot in (b) which are separated by the Relative Signal Metric and the signal to noise ratio. These plots show how the metric could be used for determining true and false detections.

IV.5.7 Removed Relative Signal Metric

Removed Relative Signal Metric is the RMS of the instantaneous relative signal residual with outliers removed. The Removed Relative Signal Metric is calculated using Equation IV.25 with relative signal residuals, but whenever the reference is close to zero, those samples are removed. This prevents the relative residual from approaching infinity as the denominator approaches zero.

Removed Relative Signal Metric is a measure of how much the estimated burst signal deviates from a reference signal. In Figure IV.14 two events with different Removed Relative Signal Metric values are shown. Figure IV.14a shows a low value for the Removed Relative Signal Metric which is indicative of an event that closely follows the reference signal. This incorporates a good estimate of amplitude, phase, and frequency to match the reference well. Figure IV.14b shows a high value for the Removed Relative Signal Metric, showing an event that does not correspond well with the reference signal indicating a poorer characterization of the event.

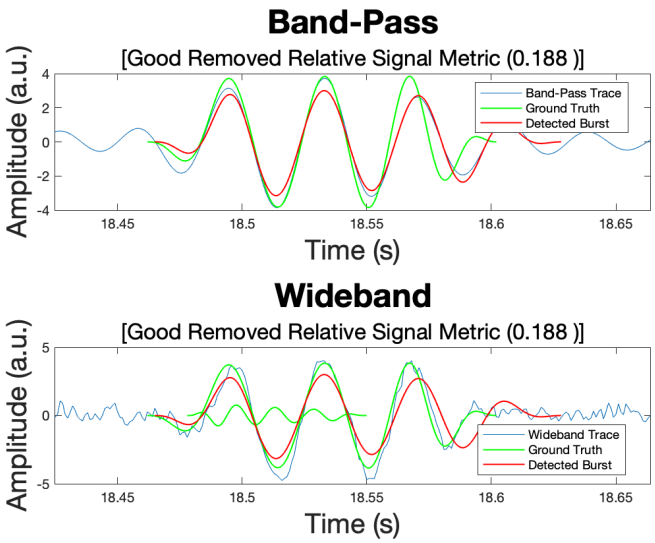

(a) Good Removed Relative Signal Metric

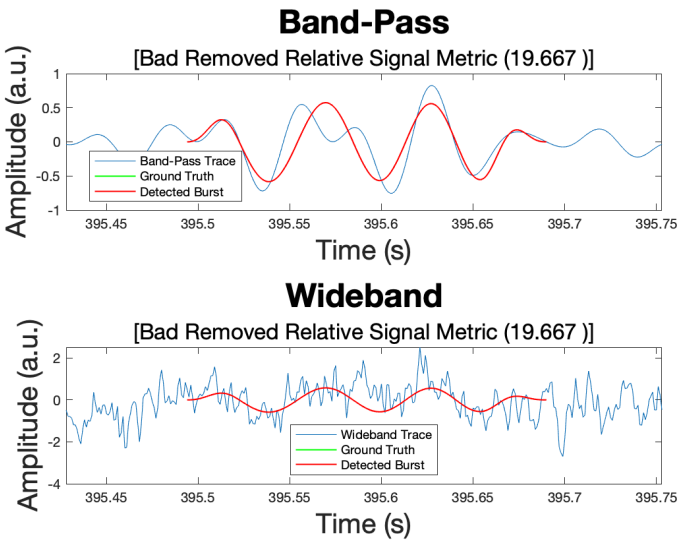

(b) Bad Removed Relative Signal Metric

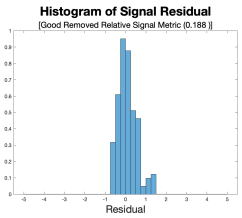

(c) Good Removed Relative Signal Metric Histogram of Residuals

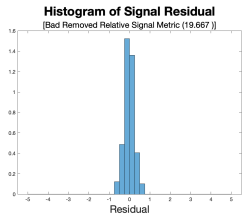

(d) Bad Removed Relative Signal Metric Histogram of Residuals

Figure IV.14: Here examples of good and bad event characterizations are shown as determined by Removed Relative Signal Metric. In (a) a low value of Removed Relative Signal Metric is shown indicating an event that closely follows the reference. In (b) a high value of Removed Relative Signal Metric is shown indicating an event that does not closely follow the reference. Both (c) and (d) show the residuals for the detected events.

Figure IV.15 shows how well this metric does at separating events by whether they represent a true or false detection. Figure IV.15a shows how the metric separates the true or false detections. The detections cluster around each other indicating that this metric cannot be used on its own to filter true detections from false detections. Figure IV.15b shows how events separates along the metric and the signal to noise ratio. While this figure shows that low SNR events are typically false positives, it does not show that the Removed Relative Signal Metric is a useful metric for separating the events by true and false detections.

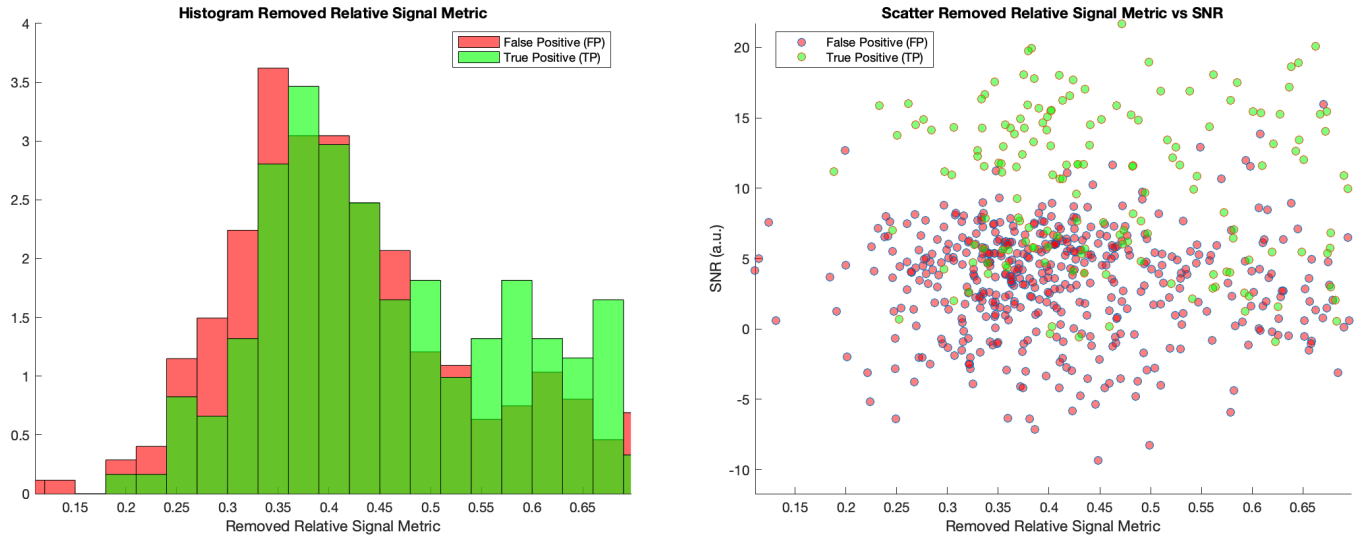

(a) Histogram Removed Relative Signal Metric

(b) Removed Relative Signal Metric Vs SNR

Figure IV.15: This figure shows the types of events, either true positive, false positive, or false negative, that the Magnitude Method generates. These events are shown in a histogram in (a) are binned by the Removed Relative Signal Metric. These events are also shown in a scatter plot in (b) which are separated by the Removed Relative Signal Metric and the signal to noise ratio. These plots show how the metric could be used for determining true and false detections.

IV.5.8 Interpolated Relative Signal Metric

Interpolated Relative Signal Metric is the RMS of the instantaneous relative signal residual with outliers removed. The Interpolated Relative Signal Metric is calculated using Equation IV.25 with relative signal residuals, but whenever the reference is close to zero, those samples are removed. This prevents the relative residual from approaching infinity as the denominator approaches zero. Additionally, after removing samples around zeros, the relative error is linearly interpolated across the removed points.

Interpolated Relative Signal Metric is a measure of how much the estimated burst signal deviates from a reference signal. In Figure IV.16 two events with different Interpolated Relative Signal Metric values are shown. Figure IV.16a shows a low value for the Interpolated Relative Signal Metric which is indicative of an event that closely follows the reference signal. This incorporates a good estimate of amplitude, phase, and frequency to match the reference well. Figure IV.16b shows a high value for the Interpolated Relative Signal Metric, showing an event that does not correspond well with the reference signal indicating a poorer characterization of the event.

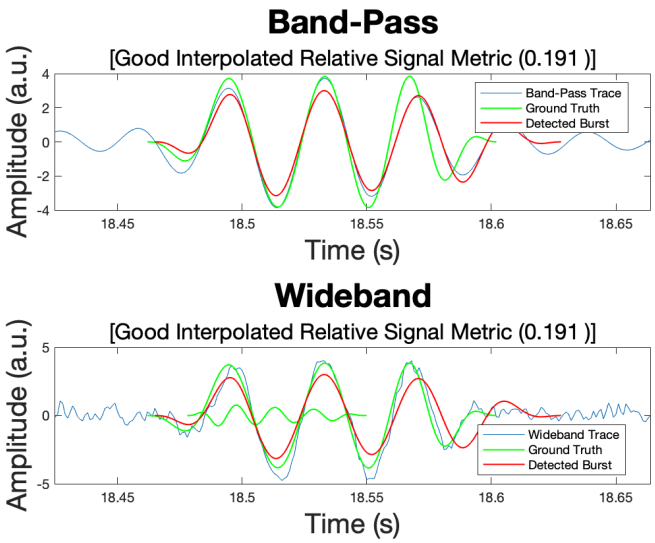

(a) Good Interpolated Relative Signal Metric

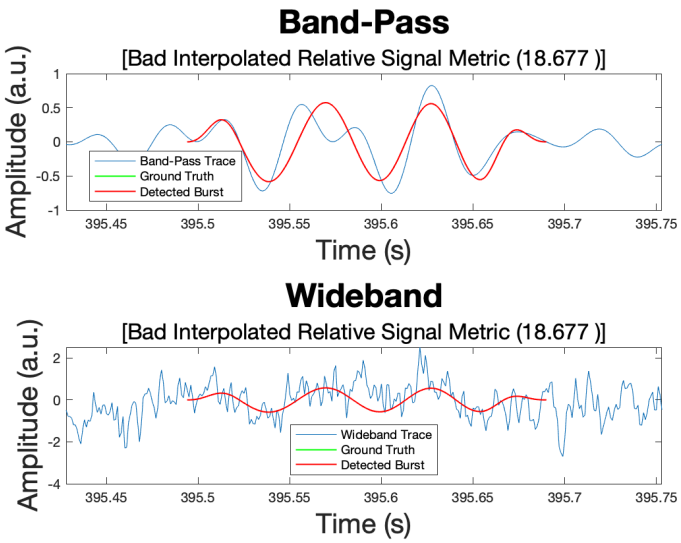

(b) Bad Interpolated Relative Signal Metric

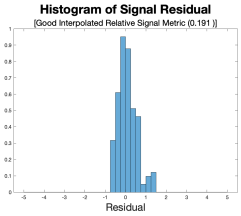

(c) Good Interpolated Relative Signal Metric Histogram of Residuals

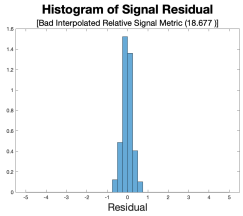

(d) Bad Interpolated Relative Signal Metric Histogram of Residuals

Figure IV.16: Here examples of good and bad event characterizations are shown as determined by Interpolated Relative Signal Metric. In (a) a low value of Interpolated Relative Signal Metric is shown indicating an event that closely follows the reference. In (b) a high value of Interpolated Relative Signal Metric is shown indicating an event that does not closely follow the reference. Both (c) and (d) show the residuals for the detected events.

Figure IV.17 shows how well this metric does at separating events by whether they represent a true or false detection. Figure IV.17a shows how the metric separates the true or false detections. The detections cluster around each other indicating that this metric cannot be used on its own to filter true detections from false detections. Figure IV.17b shows how events separates along the metric and the signal to noise ratio. While this figure shows that low SNR events are typically false positives, it does not show that the Interpolated Relative Signal Metric is a useful metric for separating the events by true and false detections. The metric performs better as a filter with a cutoff of 0.55 for events with SNR above 10.

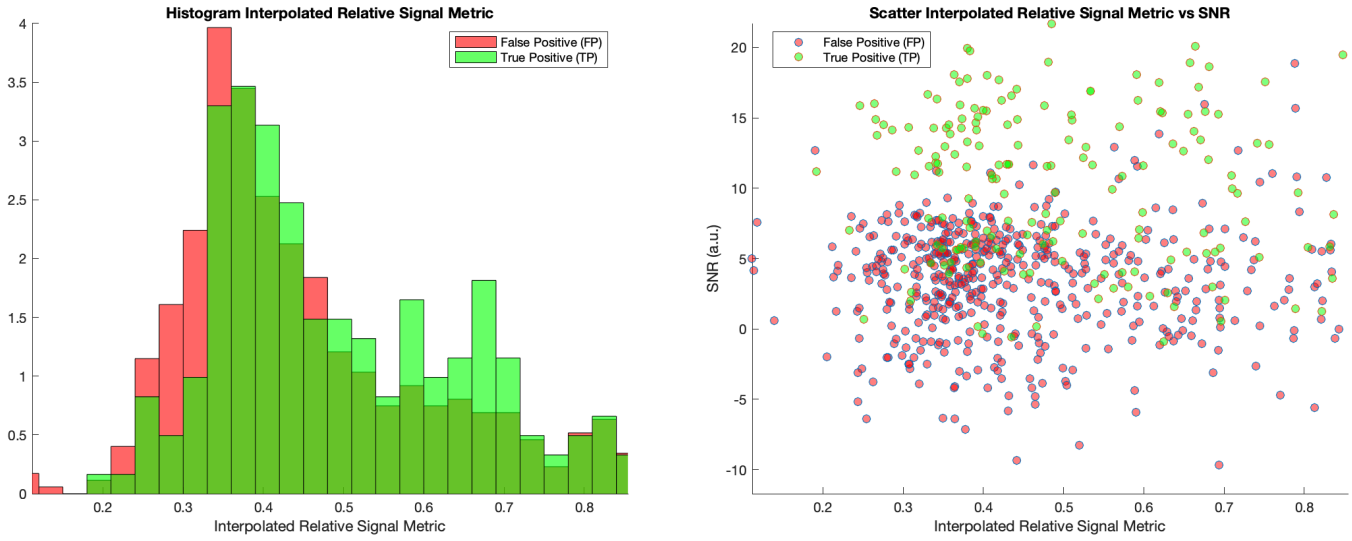

(a) Histogram Interpolated Relative Signal Metric

(b) Interpolated Relative Signal Metric Vs SNR

Figure IV.17: This figure shows the types of events, either true positive, false positive, or false negative, that the Magnitude Method generates. These events are shown in a histogram in (a) are binned by the Interpolated Relative Signal Metric. These events are also shown in a scatter plot in (b) which are separated by the Interpolated Relative Signal Metric and the signal to noise ratio. These plots show how the metric could be used for determining true and false detections.

IV.5.9 Relative Magnitude Metric

Relative Magnitude Metric is the RMS of the instantaneous relative magnitude residual. The Relative Magnitude Metric is calculated using Equation IV.25 with relative magnitude residuals, described in Equation IV.27.

The Relative Magnitude Metric is a measure of how much the estimated burst magnitude deviates from a reference magnitude. In Figure IV.18 two events with different Relative Magnitude Metric values are shown. Figure IV.18a shows a low value for the Relative Magnitude Metric which is indicative of an event that closely follows the reference magnitude. Figure IV.18b shows a high value for the Relative Magnitude Metric, showing an event that does not correspond well with the reference magnitude indicating a poorer characterization of the event magnitude.

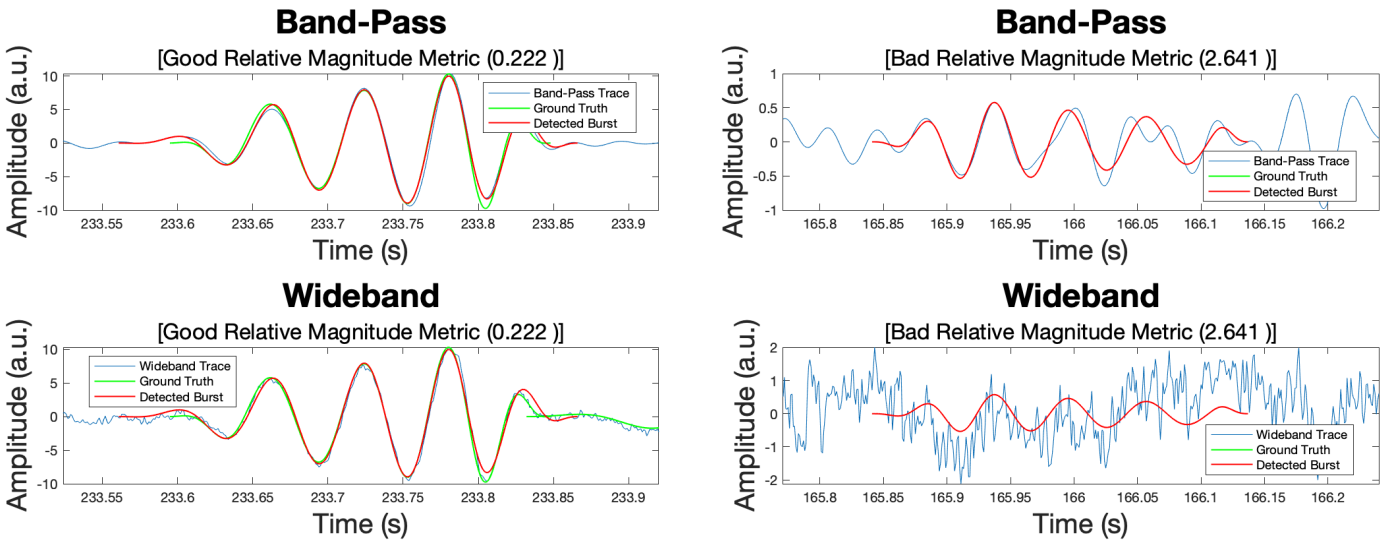

(a) Good Relative Magnitude Metric

(b) Bad Relative Magnitude Metric

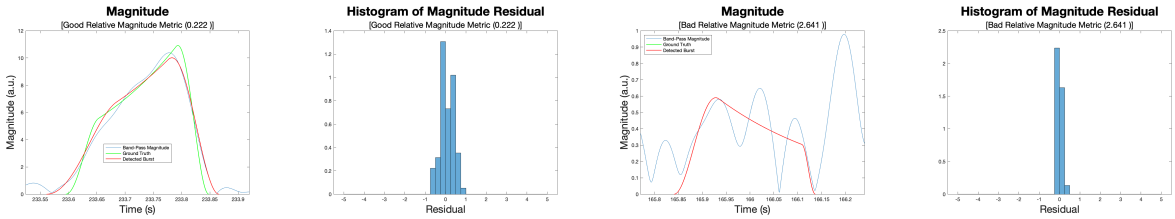

(c) Good Relative Magnitude Metric Magnitude Plot

(d) Good Relative Magnitude Metric Histogram of Residuals

(e) Bad Relative Magnitude Metric Magnitude Plot

(f) Bad Relative Magnitude Metric Histogram of Residuals

Figure IV.18: Here examples of good and bad event characterizations are shown as determined by Relative Magnitude Metric. In (a) a low value of Relative Magnitude Metric is shown indicating an event that closely follows the reference. In (b) a high value of Relative Magnitude Metric is shown indicating an event that does not closely follow the reference. Both (c) and (e) show the extracted magnitude in red and how it compares to the reference in blue, with the addition of ground truth events in green. Both (d) and (f) show the residuals for the detected events.

Figure IV.19 shows how well this metric does at separating events by whether they represent a true or false detection. Figure IV.19a shows how the metric separates the true or false detections. The detections cluster around each other indicating that this metric cannot be used on its own to filter true detections from false detections. Figure IV.19b shows how events separates along the metric and the signal to noise ratio. While this figure shows that low SNR events are typically false positives, it does not show that the Relative Magnitude Metric is a useful metric for separating the events by true and false detections. The metric performs better as a filter with a cutoff of 0.4 for events with SNR above 10.

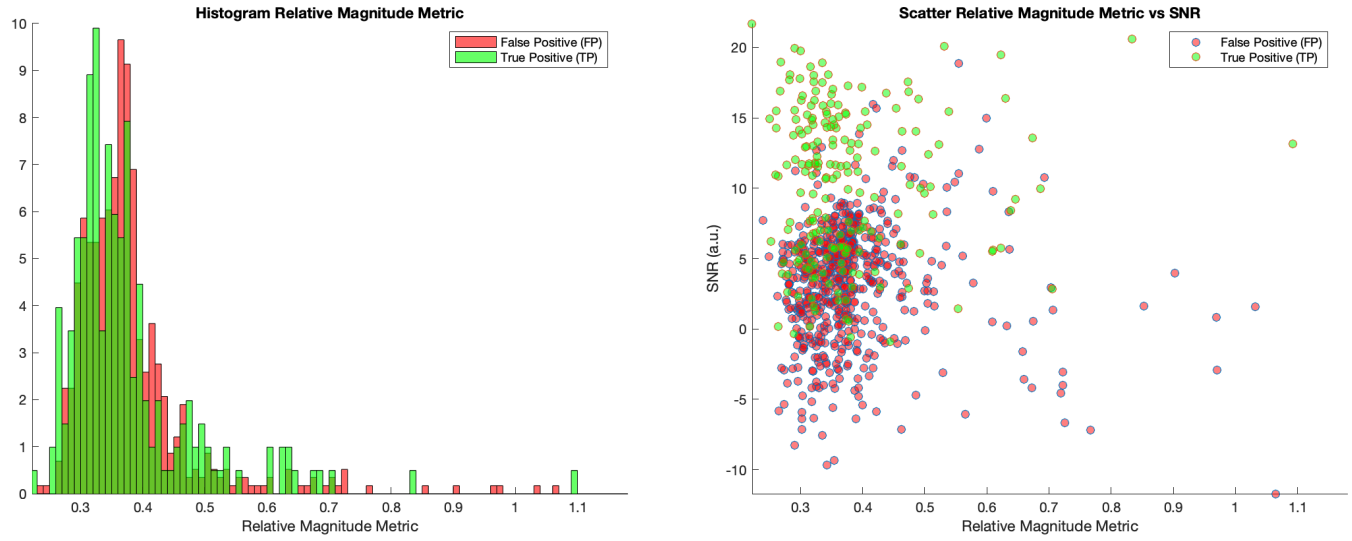

(a) Histogram Relative Magnitude Metric

(b) Relative Magnitude Metric Vs SNR

Figure IV.19: This figure shows the types of events, either true positive, false positive, or false negative, that the Magnitude Method generates. These events are shown in a histogram in (a) are binned by the Relative Magnitude Metric. These events are also shown in a scatter plot in (b) which are separated by the Relative Magnitude Metric and the signal to noise ratio. These plots show how the metric could be used for determining true and false detections.

IV.5.10 Relative Frequency Metric

Relative Frequency Metric is the RMS of the instantaneous relative frequency residual. The Relative Frequency Metric is calculated using Equation IV.25 with relative frequency residuals, described in Equation IV.27.

The Relative Frequency Metric is a measure of how much the estimated burst frequency deviates from a reference frequency. In Figure IV.20 two events with different Relative Frequency Metric values are shown. Figure IV.20a shows a low value for the Relative Frequency Metric which is indicative of an event that closely follows the reference frequency. Figure IV.20b shows a high value for the Relative Frequency Metric, showing an event that does not correspond well with the reference frequency indicating a poorer characterization of the event frequency.

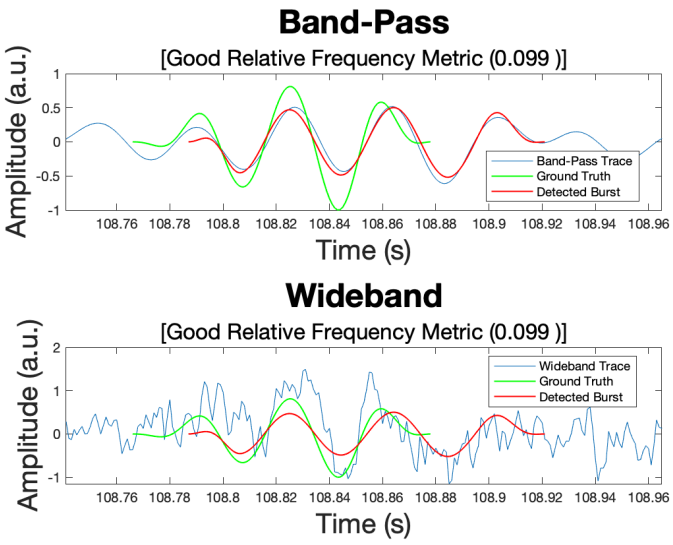

(a) Good Relative Frequency Metric

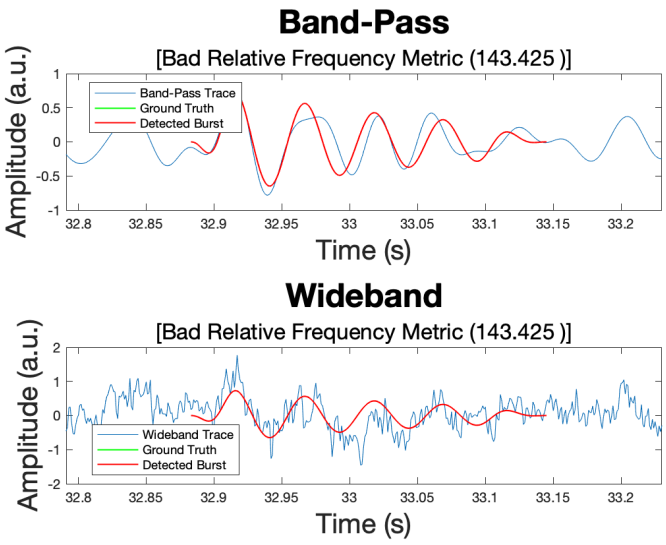

(b) Bad Relative Frequency Metric

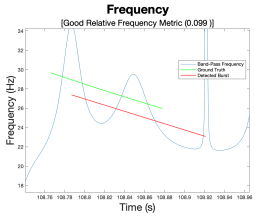

(c) Good Relative Frequency Metric Frequency Plot

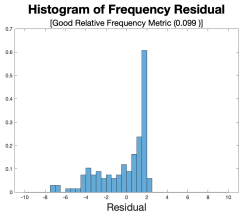

(d) Good Relative Frequency Metric Histogram of Residuals

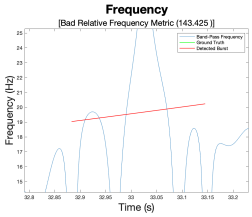

(e) Bad Relative Frequency Metric Frequency Plot

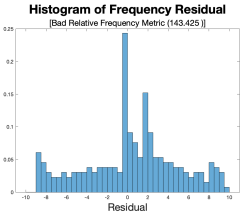

(f) Bad Relative Frequency Metric Histogram of Residuals

Figure IV.20: Here examples of good and bad event characterizations are shown as determined by Relative Frequency Metric. In (a) a low value of Relative Frequency Metric is shown indicating an event that closely follows the reference. In (b) a high value of Relative Frequency Metric is shown indicating an event that does not closely follow the reference. Both (c) and (e) show the extracted frequency in red and how it compares to the reference in blue, with the addition of ground truth events in green. Both (d) and (f) show the residuals for the detected events.

Figure IV.21 shows how this metric does at separating events by whether they represent a true or false detection. Figure IV.21a shows how the metric separates the true or false detections. The detections cluster around each other indicating that this metric cannot be used on its own to filter true detections from false detections. Figure IV.21b shows how events separates along the metric and the signal to noise ratio. While this figure shows that low SNR events are typically false positives, it does not show that the Relative Frequency Metric is a useful metric for separating the events by true and false detections.

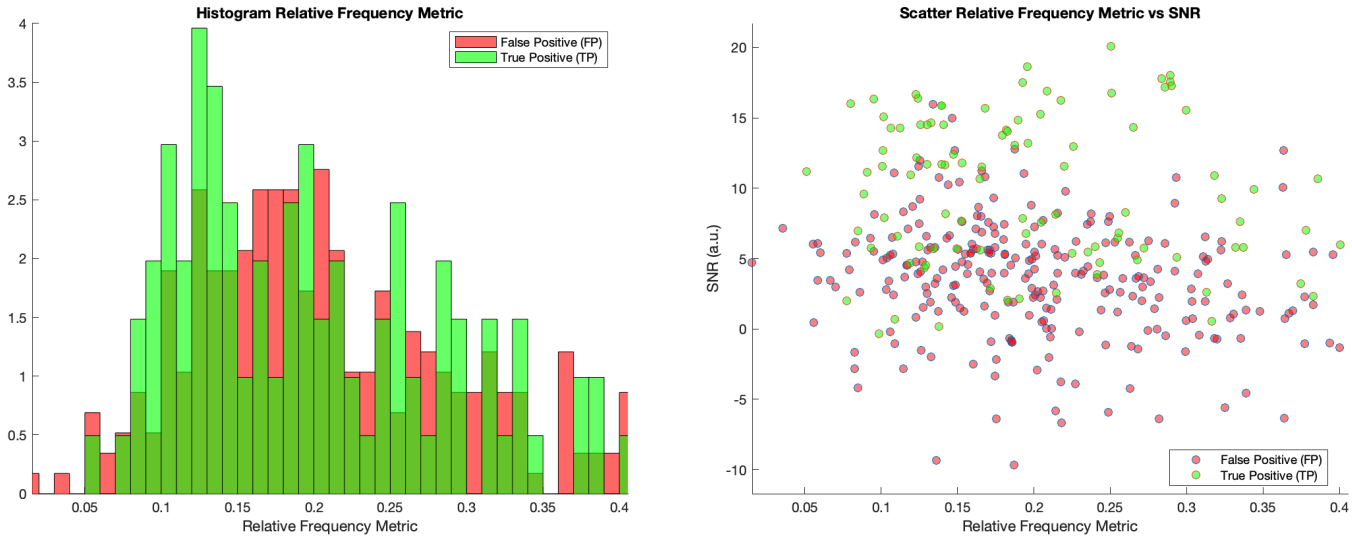

(a) Histogram Relative Frequency Metric

(b) Relative Frequency Metric Vs SNR

Figure IV.21: This figure shows the types of events, either true positive, false positive, or false negative, that the Magnitude Method generates. These events are shown in a histogram in (a) are binned by the Relative Frequency Metric. These events are also shown in a scatter plot in (b) which are separated by the Relative Frequency Metric and the signal to noise ratio. These plots show how the metric could be used for determining true and false detections.

IV.6 Circular Metrics

For metrics that relate to phase circular statistics must be used. Because phase wraps around the residuals of phase, $e_{\phi,j}$ must be treated differently. To this end three metrics are used that are derived from calculations shown in Equation IV.28. Here the phase residual, $e_{\phi,j}$, is calculated from Equation IV.24 for phase. The event is the j-th event and it has N samples.

$$\begin{aligned} C_{p,j} &= \sum_{n=1}^{N} \cos(p \cdot e_{\phi,j}(n)) \\ S_{p,j} &= \sum_{n=1}^{N} \sin(p \cdot e_{\phi,j}(n)) \\ R_{p,j} &= \sqrt{C_{p,j}^2 + S_{p,j}^2} \\ \bar{R}_{p,j} &= \frac{R_{p,j}}{N} \end{aligned} \tag{IV.28}$$

IV.6.1 Mean Direction

Mean Direction (MD) represents the average angle by which the phase estimate is off because it is an average of the phase residual which is the difference between the detected and reference phase. It can show a consistent bias in the phase estimate.

Equation IV.29 shows how it is calculated. The Mean Direction is measured relative to the phase of the band pass filtered signal. When there is a prominent ground truth event the Mean Direction should indicate the phase shift relative to the ground truth event.

$$MD_j = atan2(\bar{S}_{1,j}, \bar{C}_{1,j}) \tag{IV.29}$$

In Fig.IV.22b the detected wave has a large phase shift relative to the ground truth, causing it to have a large Mean Direction value, whereas in Fig.IV.22a the detected wave has a small phase shift relative to the ground truth, leading to a small Mean Direction value. This phase error may be averaging to a low value with Mean Direction, but could be inconsistent, which would be illustrated by a large value for Circular Variance.

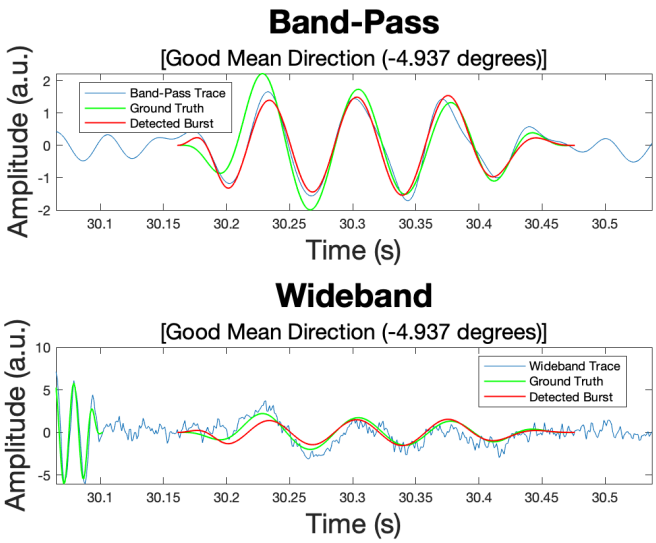

(a) Good MD

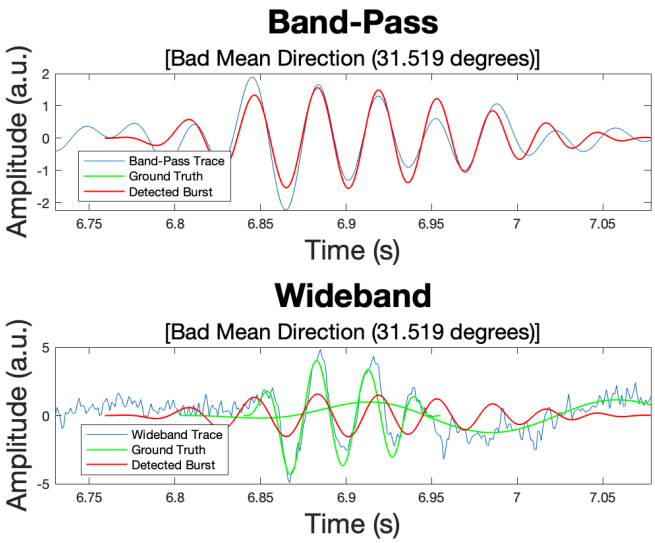

(b) Bad MD

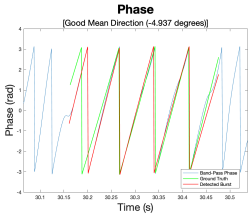

(c) Good MD Phase Plot

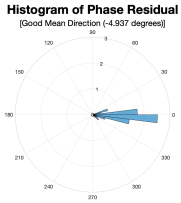

(d) Good MD Histogram of Residuals

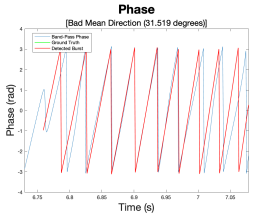

(e) Bad MD Phase Plot

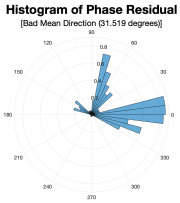

(f) Bad MD Histogram of Residuals

Figure IV.22: Here examples of good and bad event characterizations are shown as determined by MD. In (a) a low value of MD is shown indicating an event that closely follows the reference. In (b) a high value of MD is shown indicating an event that does not closely follow the reference. Both (c) and (e) show the extracted frequency in red and how it compares to the reference in blue, with the addition of ground truth events in green. Both (d) and (f) show the residuals for the detected events.

IV.6.2 Circular Variance

Circular Variance (CV) represents the random component phase error, because it is the variance of the phase residual which is the difference between the detected and reference phase. Equation IV.29 shows how it is calculated.

A circular variance between 0.8 and 1 is viewed as random whereas a value between 0 and 0.3 is viewed as autocoherent meaning that the phase is consistent [58]. A value of 0 indicates that the phase error stays constant. The Circular Variance is measured relative to the phase of the band pass filtered signal. When there is a prominent ground truth event the Circular Variance should indicate the consistency of the phase relative to the ground truth event.

$$CV_j = 1 - \bar{R}_{1,j} \tag{IV.30}$$

In Fig.IV.23b the detected wave does not consistently track the ground truth, which indicates a phase residual that is not concentrated on one value, indicating a large value for Circular Variance. However, in Fig.IV.23a the detected wave does consistently track the ground truth, which indicates a phase residual that is concentrated on one value, indicating a large value for Circular Variance. This phase error may be consistent with a low Circular Variance value, but it could be shifted, which would be illustrated by a larger Mean Direction value.

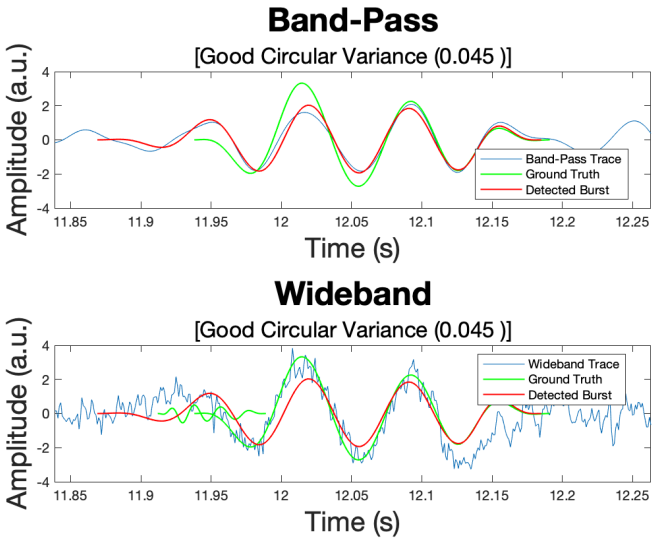

(a) Good CV

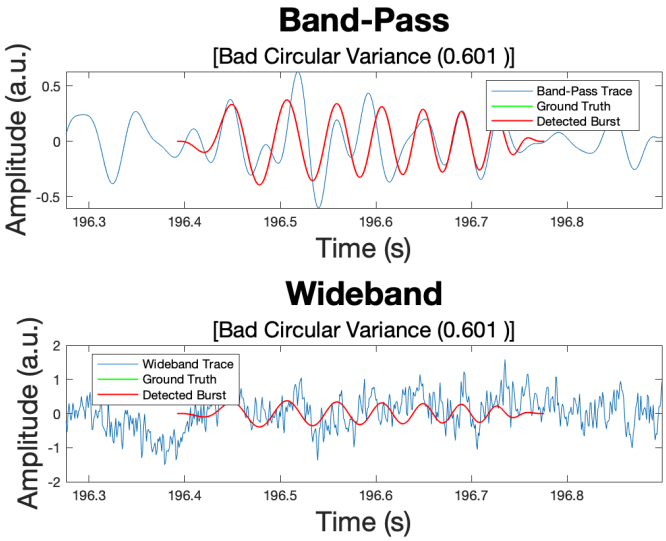

(b) Bad CV

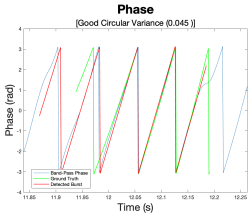

(c) Good CV Phase Plot

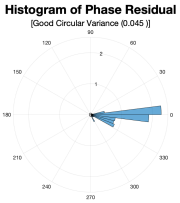

(d) Good CV Histogram of Residuals

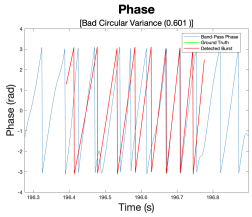

(e) Bad CV Phase Plot

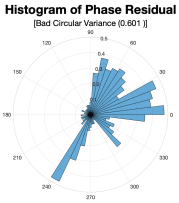

(f) Bad CV Histogram of Residuals

Figure IV.23: Here examples of good and bad event characterizations are shown as determined by CV. In (a) a low value of CV is shown indicating an event that closely follows the reference. In (b) a high value of CV is shown indicating an event that does not closely follow the reference. Both (c) and (e) show the extracted frequency in red and how it compares to the reference in blue, with the addition of ground truth events in green. Both (d) and (f) show the residuals for the detected events.

IV.6.3 Combined Angle Metric

The Combined Angle Metric (CAM) introduces a way to evaluate the phase shift and the consistency of the phase shift over the detected burst. It combines the phase shift tracking of MD with the consistency/jitter measurement of CV. Equation IV.31 shows how it is calculated.

This gives a value of 1 when the phase shift and consistency are both low indicating a good phase estimate. A lower value means that the phase shift relative to the reference is large or the phase varies greatly from the reference phase.

$$CAM_j = \sqrt{\left(1 - \frac{|MD_j|}{\pi}\right)(CV_j - 1)} \tag{IV.31}$$

In Fig.IV.24b the detected wave does not consistently track the ground truth, which indicates a phase residual that is not concentrated on one value or has a bias from the reference, leading to a large value for Combined Angle Metric. However, in Fig.IV.24a the detected wave does consistently track the ground truth, which indicates a phase residual that is concentrated on one value with a small shift from the reference, leading to a value closer to 1 for Combined Angle Metric.

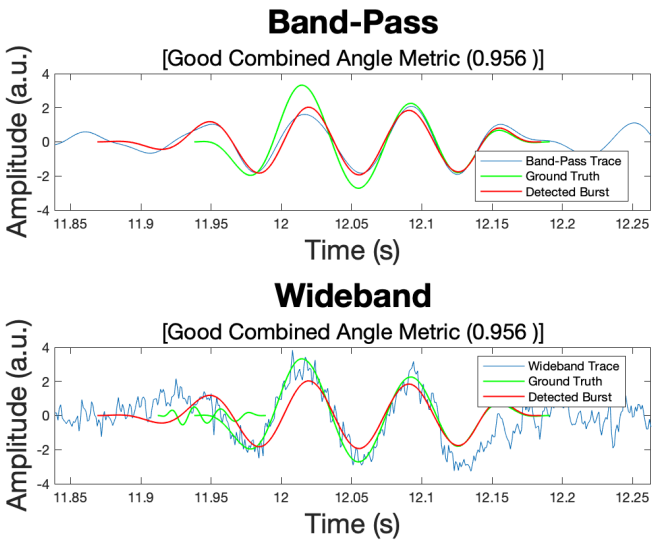

(a) Good CAM

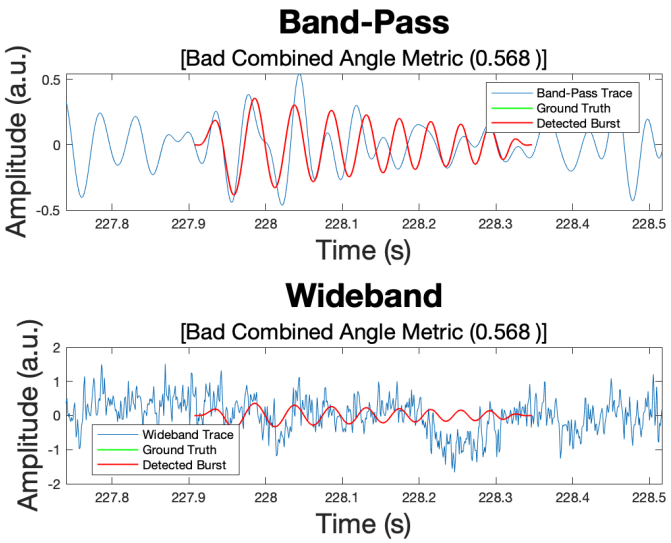

(b) Bad CAM

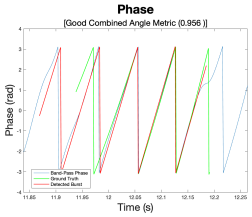

(c) Good CAM Phase Plot

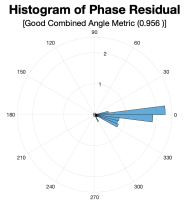

(d) Good CAM Histogram of Residuals

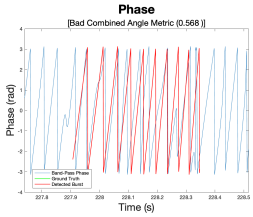

(e) Bad CAM Phase Plot

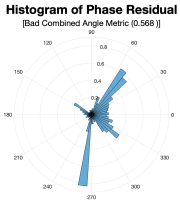

(f) Bad CAM Histogram of Residuals

Figure IV.24: Here examples of good and bad event characterizations are shown as determined by CAM. In (a) a low value of CAM is shown indicating an event that closely follows the reference. In (b) a high value of CAM is shown indicating an event that does not closely follow the reference. Both (c) and (e) show the extracted frequency in red and how it compares to the reference in blue, with the addition of ground truth events in green. Both (d) and (f) show the residuals for the detected events.

IV.7 Results

One of the aims of using metrics is to use them to tune the methods. A good metric should provide information distinct from the others, demonstrate a search space with a well-behaved optimum, and most importantly represent a feature that is desirable.

To find metrics that allow discrimination between optimal and non-optimal tuning values, and provide information distinct from other metrics they are plotted for a sweep of a primary tuning parameter, threshold, for the magnitude method. This is used to determine if the metric has the desired qualities without having to visualize over all possible tuning parameters.

Figure IV.25 shows the various metrics derived from the confusion matrix over the primary tuning parameter,

threshold, for the magnitude detection method. The primary tuning parameter, threshold, range is selected as between 0 dB and 30 dB. Any value lower than this range is below the noise floor and any value over this range does not lead to any true positive detections. In IV.25a true positive count is the metric. As the threshold is increased the number of true positive detections decreases because more ground truth events are being missed. In IV.25b false positive count is the metric. As the threshold is increased the number of false positive detections decreases because fewer events are being detected. In IV.25c false negative count is the metric. As the threshold is increased the number of false negatives increases because more ground truth events are not being detected. In IV.25d false positive rate is the metric. As the threshold is increased the false positive rate decreases meaning more of the detected events are true positive events, but this does not continue past 12 dB where is begins to increase again. In IV.25e precision is the metric. As the threshold is increased the precision increases meaning more of the detected events are true positive events, but this does not continue past 12 dB where is begins to increase again. In IV.25f recall is the metric. As the threshold is increased the recall decreases because fewer of the ground truth events are being identified correctly. In IV.25g $\mathbf{F_1}$ score is the metric. Over the range of the threshold $\mathbf{F_1}$ score has a maximum at 8 dB. In IV.25h $\mathbf{F}_\beta$ score is the metric. Over the range of the threshold $\mathbf{F}_\beta$ score has a maximum at 12 dB. Precision, $\mathbf{F_1}$ score and $\mathbf{F}_\beta$ score have distinct maxima that occur in the middle of the threshold range so they are possible candidates for tuning a method. False positive rate also has similar properties, but is not considered because it is a linear transform of precision.

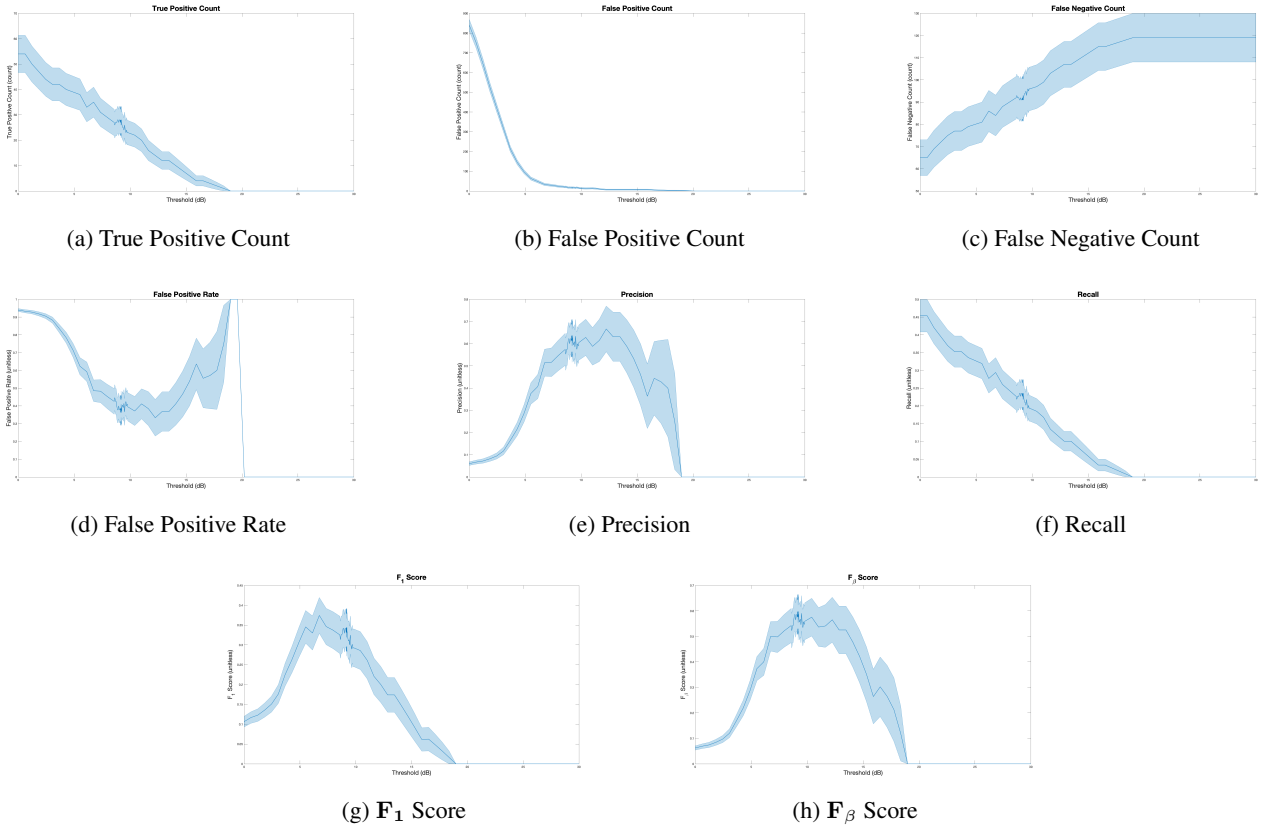

(a) True Positive Count (b) False Positive Count (c) False Negative Count

(d) False Positive Rate (e) Precision (f) Recall

(g) $\mathbf{F_1}$ Score (h) $\mathbf{F}_\beta$ Score

Figure IV.25: This figure shows the various metrics derived from the confusion matrix over the primary tuning parameter, threshold, for the Magnitude Method. To obtain these figures, the Magnitude Method was used to extract events. This was repeated for many values of the tuning parameter, threshold, to provide a view of the optimization space

Figure IV.25 it shows that $\mathbf{F_1}$ score and $\mathbf{F}_\beta$ score are good candidates for tuning. They have an optimum that is within the expected range for the primary tuning parameter that appears to be well behaved. Figure IV.25g shows a clear maximum around 8 dB, whereas Figure IV.25h shows a maximum around 12 dB. This reflects the different weighting between correctly identifying events and identifying as many true events as possible. The metrics represent a desirable measure of the detection, since a high value indicates that of the events detected more of them are true events and many of the true events are found by the detection. The other metrics do not show characteristics that are desirable for tuning.

Figure IV.26 shows the various RMS Metrics over the primary tuning parameter, threshold, for the magnitude detection method similarly to Figure IV.25. In IV.26a Signal Metric is the metric. As the threshold is increased the characterization of the events gets worse until reaching a maximum around 10 dB after which the characterization of events gets better. In IV.26b Magnitude Metric is the metric. As the threshold is increased the characterization of the magnitude of the events gets worse until reaching a maximum around 10 dB after which the characterization

of events gets better. In IV.26c Frequency Metric is the metric. As the threshold is increased the characterization of the frequency of the events improves but the improvement is modest until no events are detected. In IV.26d Relative Event Power to Wideband Metric is the metric. As the threshold is increased the characterization improves until reaching a minimum around 14 dB after which the characterization of events gets worse. In IV.26e Relative Event Power to Bandpass Metric is the metric. As the threshold is increased the characterization of events continually worsens until no more events are detected. In IV.26f Relative Frequency Metric is the metric. As the threshold is increased the characterization of the frequency of the events improves until no more events are detected. In IV.26g Relative Magnitude Metric is the metric. As the threshold is increased the characterization of the magnitude of events worsens until no more events are detected. In IV.26h Relative Signal Metric is the metric. As the threshold is increased the characterization of the events worsens until no more events are detected with an increase at the start due to edge effects. In IV.26i Removed Relative Signal Metric is the metric. As the threshold is increased the characterization of the events worsens until no more events are detected. In IV.26j Interpolated Relative Signal Metric is the metric. As the threshold is increased the characterization of the events worsens until no more events are detected.

Figure IV.26 it shows that Relative Event Power to Wideband Metric might be a good candidate for tuning. It have an optimum that is within the expected range for the primary tuning parameter that appears to be well behaved. Figure IV.26d shows a clear minimum around 14 dB. Additionally, a minimum for Relative Event Power to Wideband Metric represents a better characterization of events, as opposed to a maximum, which would represent worsening characterization such as with Signal Metric. However, Relative Event Power to Wideband Metric does not quite represent a value that would be useful for experiments since it does not incorporate accurate burst detection or phase characterization.

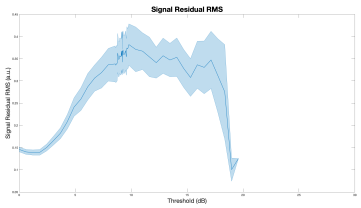

(a) Signal Metric

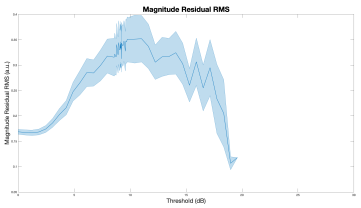

(b) Magnitude Metric

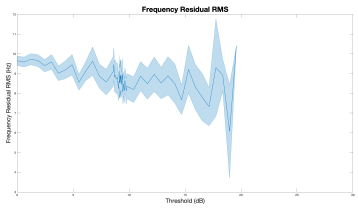

(c) Frequency Metric

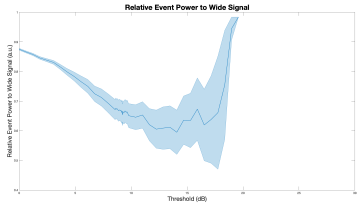

(d) Relative Event Power to Wideband Metric

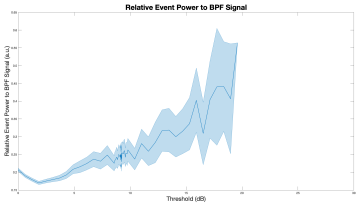

(e) Relative Event Power to Bandpass Metric

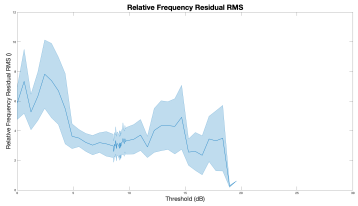

(f) Relative Frequency Metric

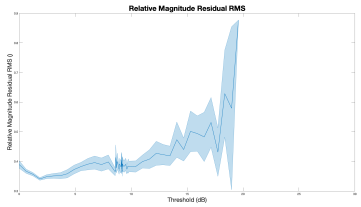

(g) Relative Magnitude Metric

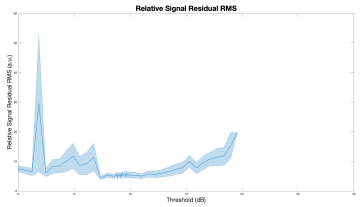

(h) Relative Signal Metric

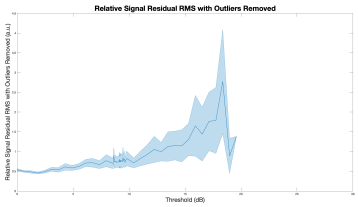

(i) Removed Relative Signal Metric

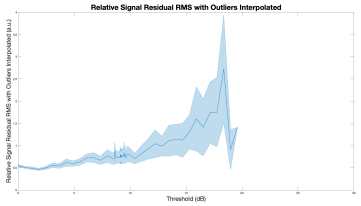

(j) Interpolated Relative Signal Metric

Figure IV.26: This figure shows the various RMS Metrics over the primary tuning parameter, threshold, for the Magnitude Method. To obtain these figures, the Magnitude Method was used to extract events. This was repeated for many values of the tuning parameter, threshold, to provide a view of the optimization space

Figure IV.27 shows the various Circular Metrics over the primary tuning parameter, threshold, for the magnitude detection method similarly to Figure IV.25. In IV.27a CV is the metric. As the threshold is increased the characterization of the phase of the events becomes more varied until no events are detected. In IV.27b MD is the metric. As the threshold is increased the characterization of the phase of the events stays fairly constant until worsening at the end. In IV.27c CAM is the metric. As the threshold is increased the characterization of the phase of the events becomes worse until no events are detected.

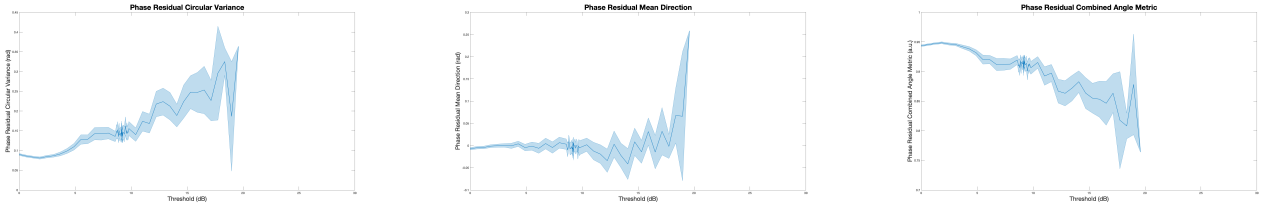

(a) Circular Variance (b) Mean Direction (c) Combined Angle Metric

Figure IV.27: This figure shows the various Circular Metrics over the primary tuning parameter, threshold, for the Magnitude Method. To obtain these figures, the Magnitude Method was used to extract events. This was repeated for many values of the tuning parameter, threshold, to provide a view of the optimization space

Figure IV.27 does not show any metrics that would be good candidates for tuning based on the behavior for the primary tuning parameter because none of them have clear optimum. However, CV does represent a value that would be useful for experiments since it incorporates one of two desirable aspects which is phase characterization. If CV showed a clear maximum then it would be a good candidate for tuning methods.

IV.8 Conclusion

There are several families of performance metrics described, confusion matrix based metrics, circular statistic based metrics, and RMS based metrics. These different families allow for a more complete comparison across burst detection and characterization methods.

Characterization metrics show some ability to discriminate between True Positive and False Positive events. Future work will be to use these as post-detection filters or clustering for detection algorithms, which the `wlBurst_v2` library already does with the frequency detection.

Additionally, these metrics enable tuning. The primary tuning parameter figures show how the metrics behave as an objective function. Only $\mathbf{F_1}$ score, $\mathbf{F}_\beta$ score, and Relative Event Power to Wideband Metric appeared to be well-behaved objective functions. While these could be used as objective functions for tuning they did not all represent desirable metrics for the phase specific experiments. $\mathbf{F_1}$ score, $\mathbf{F}_\beta$ score, and CAM do represent desirable metrics since they measure accurate detection of bursts or accurate characterization of phase.

Given this information, the $\mathbf{F}_\beta$ score will be used for tuning. It provides a desirable metric and shows good behavior for possible tuning.

CHAPTER V

Metric-Guided Tuning of Detection and Characterization Algorithms

V.1 Introduction

The methods that were described in chapter III have many tuning parameters that affect the outcome of the detection and characterization. Parameters that are expected to be tuned for a given experiment are designated as primary tuning parameters. One example is the magnitude threshold for the Hilbert Magnitude method. Other parameters that are expected to be set during the design of the algorithm are designated as secondary tuning parameters. An example of a secondary tuning parameter is the dropout parameter, which is used to ignore times when the magnitude briefly falls below-threshold in the `wlBurst_v2` implementation of the Hilbert Magnitude method.

Setting both types of parameter is usually ad-hoc. For primary parameters, traditional values often exist such as 2-σ or 3-σ for magnitude thresholds [1] or 98th percentile for wavelet thresholds [72]. For secondary parameters, the choice is typically arbitrary and/or based on trial and error. The purpose of this experiment is to apply automated optimization approaches to tuning primary and secondary parameters, so that these may be set in a systematic way and may provide performance that is closer to optimal, for some given choice of performance measurement.

The goal is to find a global maximum of one performance metric across tuning parameter space. For this optimization problem the objective function is the metric value obtained after detection and characterization of the data. Constraints using additional performance metrics may be explicitly implemented or may be considered modifications of the objective function such as dropping the value to zero if another specified metric goes outside the permitted range.

Baseline global search method is a brute force grid search, offered by `wlBurst_v2`, but is expensive for two parameters and prohibitive for more than two. Many local search, hill-climbing, approaches exist that analyze a local region of parameter space around a point and incrementally move that point towards a local maximum which imay or may not be the global maximum [62, 65, 68]. For a smooth well-behaved objective function this is a plausible method, but the objective function may include many local maxima especially if it is discontinuous or sampled.

Two approaches are used in this work: (1) an iterated univariate grid search sometimes called one-at-a-time optimization [64, 65]. It globally searches each axis of the tuning parameter space in succession. (2) A random search algorithm sometimes called creeping random search from the `wlBurst_v2` library [61–63, 65–69]. It explores neighborhoods of points with some probability distribution around a starting point and moving to any better point found.

Chapter IV showed that the $\mathbf{F}_\beta$ score is a good metric to use for tuning. For all of the tuning the $\mathbf{F}_\beta$ score will be used.

The following sections describe optimization approaches, describe the choice of metric being optimized, list the algorithms used for the test, and describe the primary and secondary tuning parameters for these algorithms

V.2 Methods

Here the optimization experiments are described. First is an example of the data structure that describes the tuning parameters. Next is a description of the grid search method and the stochastic gradient descent method. Examples are given to show how these two methods work on a two dimensional problem.

Table V.1 shows an example data structure used for optimization. This is the data structure for the Wavelet Method. The first field is the parameter name. This includes parameters that are never tuned such as the detection function since they determine the method to use. It also inclused primary tuning parameters such as threshold, and secondary tuning parameters such as support sizes. The next column, Parameter Type, indicates whether the parameter is primary, secondary, or neither. Primary tuning parameters are those that are typically tuned and secondary are those that are normally set ad-hoc. The next column indicates the range of the parameter. Method Specific is used for parameters that do not have a range because they specify the method to use. Real is used for parameters that can take any real value. Binary is used for a parameter that can have only two values. Lastly Integer is used for parameters that can take on only integer values. The structure also includes field for a boolean for whether the tuning parameter should be tuned. A value of 1 indicates that this parameter will be tuned in the optimization, whereas 0 means that it will stay constant for the optimization. Lastly there are fields for maximum and minimum values of the tuning parameter. The first value represents the theoretical maximum or minimum and the second value represents the practical maximum and minimum. This structure is used to guide the optimizations.

Parameter Name	Parameter Type	Range Type	Want Tune	Maximum Value	Minimum Value
Type of Detection Function	0	Method Specific	0	[NaN,NaN]	[NaN,NaN]
Detection Function	0	Method Specific	0	[NaN,NaN]	[NaN,NaN]
Burst Threshold (dB)	1	Real	1	[Inf,30]	[-Inf,0]
Maximum Frequency (Hz)	2	Real	1	[Inf,100]	[1,1]
Time Close Support Size (period)	2	Real	1	[Inf,10]	[0,0]
Frequency Close Support Size (decade)	2	Real	1	[Inf,15]	[0,0]
Time Open Support Size (period)	2	Real	1	[Inf,10]	[0,0]
Frequency Open Support Size (decade)	2	Real	1	[Inf,15]	[0,0]
Border Connectivity	2	Binary	1	[8,8]	[4,4]
Use Constant Q Transform	0	Binary	0	[1,1]	[0,0]
Bins Per Octave	0	Integer	0	[96,96]	[1,1]
Record All Parameters	0	Binary	0	[1,1]	[0,0]
Background Window Size (s)	2	Real	1	[Inf,20]	[0,0]
Background Window Stride (s)	2	Real	1	[Inf,20]	[0,0]
Use Magnitude for Calculations	0	Binary	0	[1,1]	[0,0]

Table V.1: This shows an example of the data structure used to describe the tuning parameters to be used in the various optimization methods for the Wavelet Method. The first column is the parameter name. Each adjustable component of the method is listed. The next column, Parameter Type, describes whether the parameter is typically tuned (primary) denoted with a 1, able to be tuned but not typically (secondary) denoted with a 2or not able to be tuned (neither) denoted with a 0. The next column is range type. These can be Method Specific, Real, Binary, or Integer. Method Specific is used for parameters that should not vary. Real is used for parameters that can vary over the real numbers. Integer is used for parameters that can only vary over integers. Binary is used for parameters that can only vary over two values. The next column describes a Boolean for whether the parameter should be altered for the optimization process. If the value is 1 then this parameter will be tuned during the optimization. If the value is 0 then this parameter will be kept constant over the optimization. The last two columns are the maximum and minimum values that the parameters can assume. There are two values in each column. The first value represents the theoretical maximum or minimum, while the second value represents the practical maximum or minimum.

V.2.1 Grid Search

The grid search first selects a tuning parameter then divides the suitable range of the parameter into linearly spaced points depending on the number of probes selected. Evaluation of the objective function is performed with each value of the tuning parameter set by the spacing with all other tuning parameters held constant. The optimal value for the objective function is used to determine the tuning parameter value. The range is then changed to the value above and below the selected tuning parameter value and this range is divided into linearly spaced points depending on the number of probes selected. The number of levels of fine tuning is a parameter of the optimization. After a value of the tuning parameter is chosen the next tuning parameter is chosen in the same manner.

Algorithm 6 Grid Search Algorithm

Initialize $\mathbf{x}^* = \mathbf{x}_0$ where x_0 is the initial values for each parameter and $\mathbf{x}^*$ is the optimum
for $i = 1$ to Number of Loops **do**
  for $j = 1$ to Number of Tuning Parameters **do**
    Where the number of tuning parameters is equal to the size of $\mathbf{x}_0$ and x_j represents the j-th value in $\mathbf{x}$
    for $k = 1$ to Number of Granularity Levels **do**
      $S_x = \emptyset$ where S_x is the set of values for the j-th parameter for the current Granularity level
      $S_f = \emptyset$
      Select Number of Probes, N, for this Level
      if $k = 1$ First Level **then**
        Select maximum and minimum for x_j from range of the input signal
      else
        Select maximum and minimum for x_j from $x_j \pm \Delta x_j$
      end if
      Calculate current Range, r, from maximum and minimum
      $\Delta x_j = \frac{r}{N}$
      Construct set of test points S_x from the maximum to the minimum at intervals of Δx_j, for replacing x_j in $\mathbf{x}^*$
      for $l = 1$ to Number of Probes **do**
        $\mathbf{x}' = \mathbf{x}^*$
        Replace the j-th parameter of $\mathbf{x}'$ with x_j selected from S_x
        if $f(\mathbf{x}') > f(\mathbf{x}^*)$ where f is the objective function **then**
          $\mathbf{x}^* = \mathbf{x}'$
        end if
      end for
    end for
  end for
end for
return $\mathbf{x}^*$ and $f(\mathbf{x}^*)$

V.2.2 Creeping Random Search

The creeping random search method works by selecting random displacements from an original point and generating a new starting point, only if the value of the objective function is better. First a random vector is sampled from a multivariate normal distribution with zero mean and unit standard deviation,which generates uniform distribution of the direction. This vector is scaled by a random value uniformly distributed in the log domain, allowing for insensitivity to characteristic scales in the search space at a small performance penalty. Lastly, this vector is scaled by a scale factor for each dimension. This displacement vector is added the original point. The value of the objective function is computed and if the value is greater than the previous value then the starting point is updated along with the new best value.

Algorithm 7 Creeping Random Search Algorithm

1: Initialize $\mathbf{x}^* = \mathbf{x}_0$ where $\mathbf{x}_0 \in \mathbb{R}^n$
2: Initialize $N_f = 0$ and $N_t = 0$ where N_f is the number of failed probes and N_t is the total number of probes
3: **while** $N_f$is less than the number of allowable failed probes and N_t is less than the number of allowable total probes
4: Create a random vector $\mathbf{x}^+ \in \mathbb{R}^n$
5: Scale $\mathbf{x}^+$ by scale vector $\mathbf{s} \in \mathbb{R}^n$ with elementwise multiplication obtaining ${\mathbf{x}_\mathbf{s}}^+$
6: $\mathbf{x}' = \mathbf{x}^* + {\mathbf{x}_\mathbf{s}}^+$
7: **if** $f(\mathbf{x}') > f(\mathbf{x}^*)$ **then**
8: $\mathbf{x}^* = \mathbf{x}'$
9: $N_f = 0$ and $N_f = N_f + 1$
10: **else**
11: $N_f = N_f + 1$ and $N_f = N_f + 1$
12: **end if**
13: **return** $\mathbf{x}^*$ and $f(\mathbf{x}^*)$

V.2.3 Optimization Examples

In this part the optimization methods are tested on three different search spaces: a smooth Gaussian, two Gaussians with different maxima, and single Gaussian distribution with low passed Gaussian noise. The termination criteria of these optimizations was that a predetermined number of iterations was reached. The optimization was determined to have reached a maximum if it identifies a maximum in the area of the known maximum.

Figure V.1 shows examples of a smooth Gaussian distribution. These show how each optimization may perform under an ideal condition. In Figure V.1a the grid search optimization climbs the hill along a diagonal, whereas in Figure V.1b the creeping random search optimization climbs more directly to the maximum.

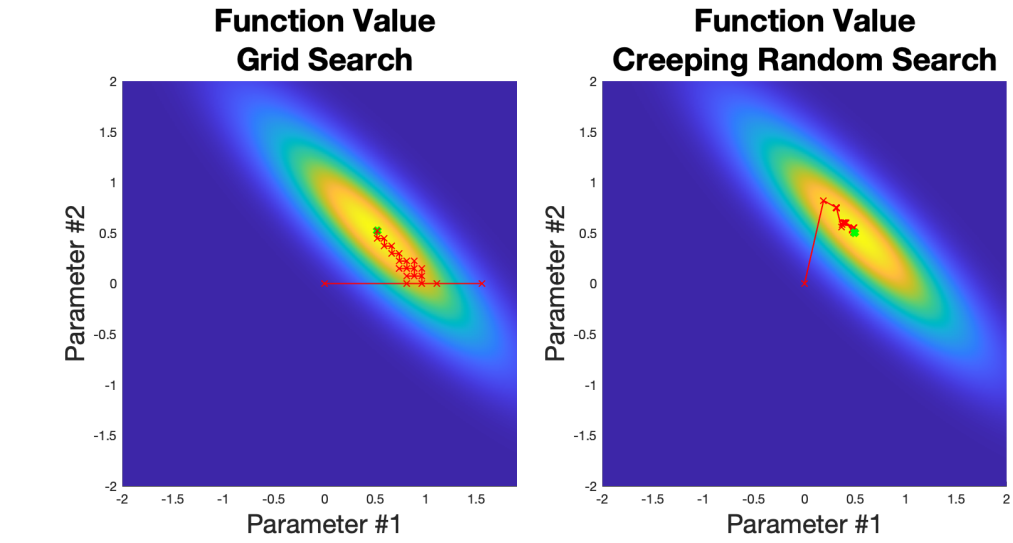

(a) Grid Search (b) Creeping Random Search

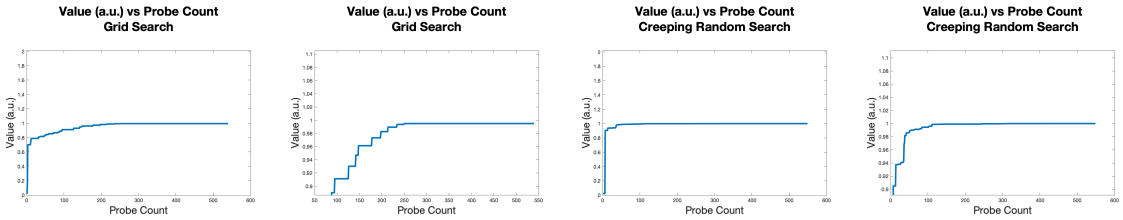

(c) Grid Search (d) Grid Search (Zoomed) (e) Creeping Random Search (f) Creeping Random Search (Zoomed)

Figure V.1: Example function used to illustrate how the grid search and creeping random search perform on a simple optimization. In (a) the grid search method is shown reaching the maximum by searching for maxima along each direction. In (b) the creeping random search method is shown reaching the maximum after several iterations. The function value per iteration for the grid search method is shown in (c) and (d). Both show the quick incremental increase in the identified maximum. In (e) and (f) the standard and zoomed iteration of the creeping random search are shown respectively. They show that the creeping random search method locates the maximum as well.

Figure V.2 shows examples of two Gaussians with different maxima. These show how each optimization may perform under conditions where two maxima exist. In Figure V.2a the grid search optimization climbs to a maximum, but it is not the global maximum, whereas in Figure V.2b the Creeping Random Search optimization climbs to the closer maxima, but tunnels to the other hill to find the global maximum. This example shows how the grid search method may fail if there are local maxima that are not orthogonal to the global maximum, but the Creeping Random Search method is able to visit the local maxima, but keep moving to the global maxima.

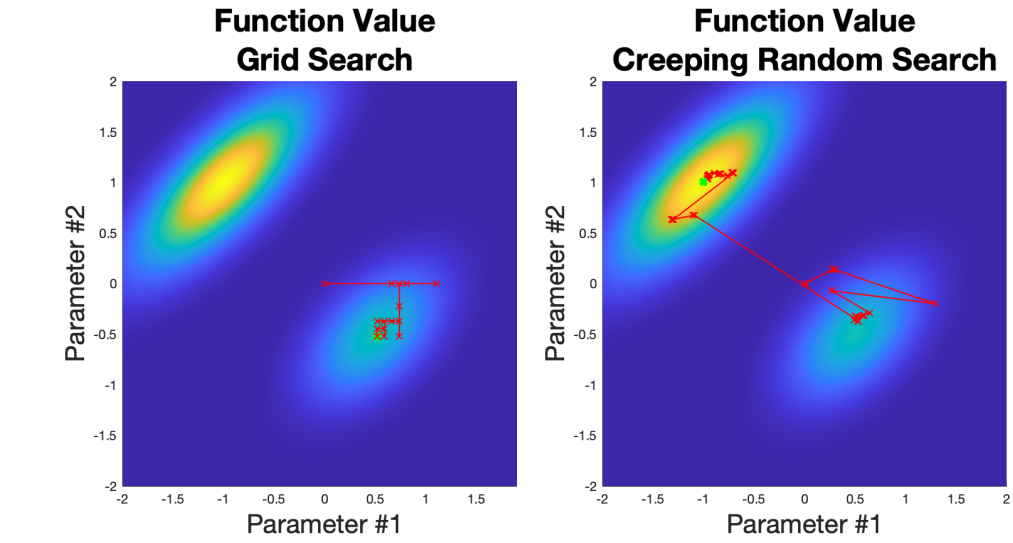

(a) Grid Search (b) Creeping Random Search

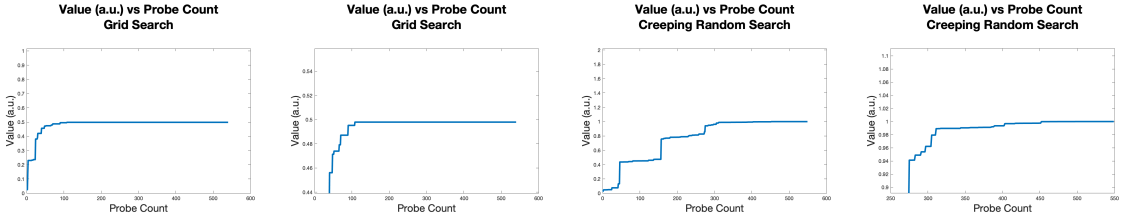

(c) Grid Search (d) Grid Search (Zoomed) (e) Creeping Random Search (f) Creeping Random Search (Zoomed)

Figure V.2: Example function used to illustrate how the grid search optimization may get stuck where the Creeping Random Search optimization will find the global optimum. In (a) the grid search method is shown starting close to a local maximum and finding that instead of the global maximum. In (b) the Creeping Random Search method is shown starting close to a local maximum, but finding the global maximum. The function value per iteration for the grid search method is shown in (c) and (d). Both show the quick location of the local maximum. In (e) and (f) the standard and zoomed iteration of the Creeping Random Search are shown respectively. They show that the Creeping Random Search method locates the local maximum, but then finds the global maximum.

Figure V.3 shows examples of a single Gaussian distribution with low passed Gaussian noise. This noise simulates many possible maxima. These show how each optimization may perform under conditions where multiple maxima exist and noise heavily distorts the search path. In Figure V.3a the grid search optimization climbs to a maximum, but it is not the noiseless maximum as previously seen, but rather a maximum created by the noise. In Figure V.3b the Creeping Random Search optimization climbs to the global maximum, but passes by other maxima created by the noise. This example shows how the grid search method may fail if the noise is such that it cannot find maxima in any orthogonal direction, but the Creeping Random Search method is able to visit the maxima, but keep moving to the global maxima.

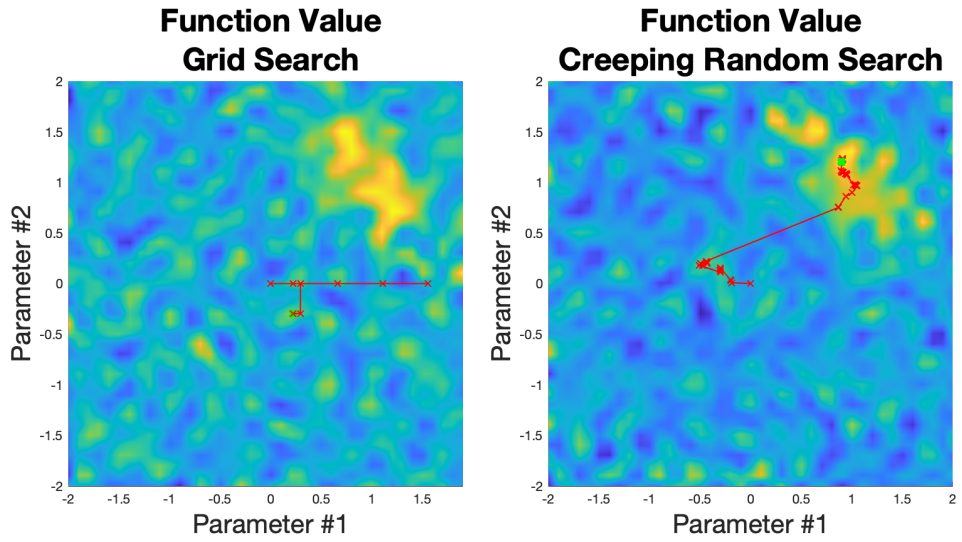

(a) Grid Search (b) Creeping Random Search

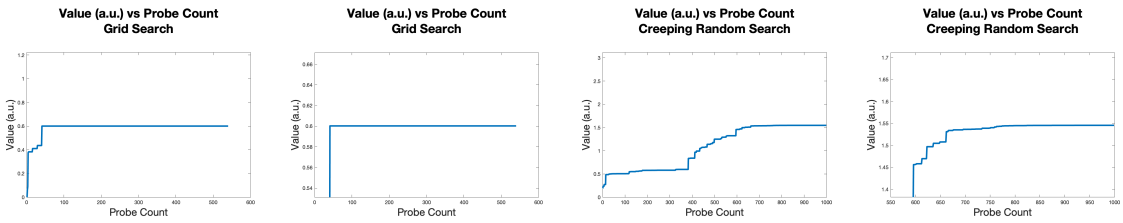

(c) Grid Search (d) Grid Search (Zoomed) (e) Creeping Random Search (f) Creeping Random Search (Zoomed)

Figure V.3: Example function used to illustrate how the grid search optimization may get stuck where the Creeping Random Search optimization will find the global optimum when there are many local optima. The objective function was constructed from a Gaussian with low frequency noise to mimic several optima. In (a) the grid search method is shown finding a local maximum that is close to the starting point, but does not represent the global maximum. In (b) the Creeping Random Search method is shown starting at the same point, but finding the global maximum. It finds several local maxima, but eventually converges on the global maxima. The function value per iteration for the grid search method is shown in (c) and (d). Both show the quick location of a local maximum. In (e) and (f) the standard and zoomed iteration of the Creeping Random Search are shown respectively. They show that the Creeping Random Search method locates several local maxima, but then finds the global maximum.

V.3 Results

To tune the methods described in chapter III an objective function is necessary. In chapter IV many metrics are explored and these can be used as objective functions. A good objective function for tuning the methods should provide information distinct from the others, demonstrate a search space with a well-behaved optimum, and most importantly represent a feature that is desirable. As discussed in chapter IV, the $\mathbf{F_1}$ score or $\mathbf{F}_\beta$ scores are best suited. They represent a desirable measure of the detection, since a high value indicates that of the events detected more of them are true events and many of the true events are found by the detection. While they do not measure the performance of the burst characterization, they are still good candidates for objective functions.

Here $\mathbf{F}_\beta$ score is chosen with $\beta = 0.2$, because it has been used with neural data before with the given value of β [33]. Tuning using both optimization methods was performed on both synthetic and real data. The initial parameter settings were obtained from literature where available or set ad-hoc if there were not exact matches to literature values.

The tuning for $\mathbf{F}_\beta$ score on synthetic data is shown in Figure V.4. This shows a comparison of the methods and how they perform on $\mathbf{F}_\beta$ Score. It also shows the improvement that is made by tuning the primary or all parameters. For the Frequency and Peak and Trough methods creeping random search does not improve detection when all parameters are allowed to vary. This is because there was not enough iterations to allow for convergence since the iteration count was the stopping criteria for optimization.

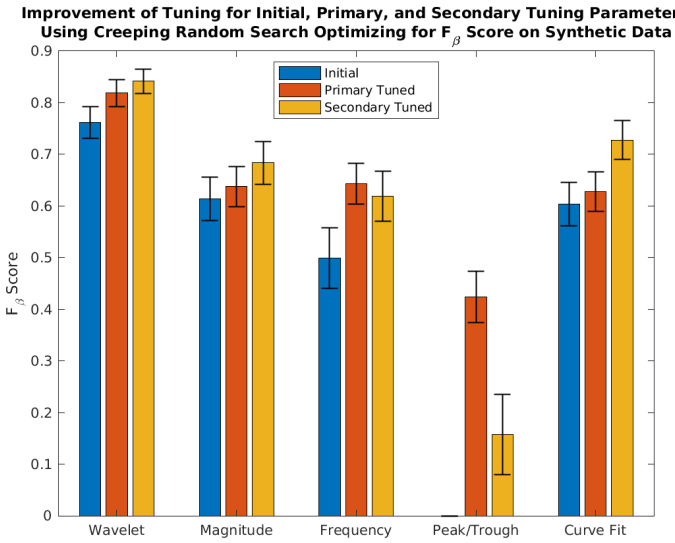

Figure V.4: Tuning results for given methods. The methods were tuned using $\mathbf{F}_\beta$ Score, with $\beta = 0.2$, on synthetic data using creeping random search for optimization. The initial point is given by the standard parameter values for the given method. The primary tuning is the result of tuning for $\mathbf{F}_\beta$ score and only adjusting the primary tuning parameter that would normally be adjusted when a method is used. The secondary tuning is the result of tuning for $\mathbf{F}_\beta$ Score and adjusting all adjustable parameters of the method. This chart shows that tuning primary parameters will lead to a more desired method and that tuning for secondary tuning parameters leads to even more of a desired result measured with $\mathbf{F}_\beta$ score.

The tuning for $\mathbf{F}_\beta$ score on synthetic data is shown in Figure V.5. This shows a comparison of the methods and how they perform on $\mathbf{F}_\beta$ Score. It also shows the improvement that is made by tuning the primary or all parameters. All methods improved when more parameters were tuned.

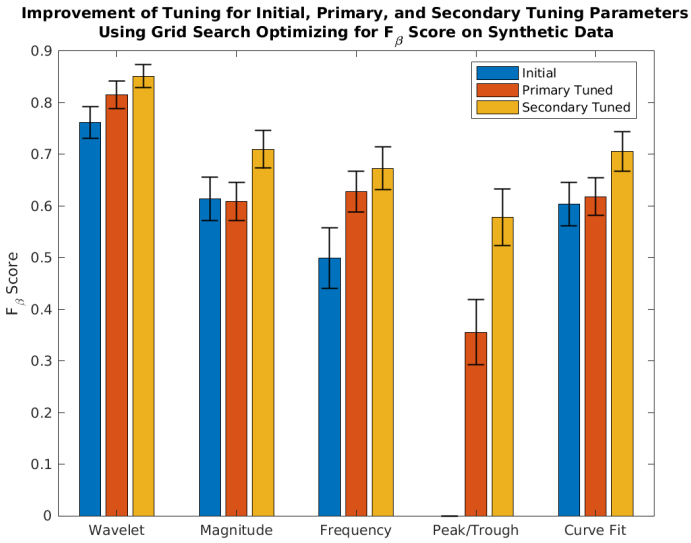

Figure V.5: Tuning results for given methods. The methods were tuned using $\mathbf{F}_\beta$ Score, with $\beta = 0.2$, on synthetic data using univariate grid search for optimization. The initial point is given by the standard parameter values for the given method. The primary tuning is the result of tuning for $\mathbf{F}_\beta$ score and only adjusting the primary tuning parameter that would normally be adjusted when a method is used. The secondary tuning is the result of tuning for $\mathbf{F}_\beta$ Score and adjusting all adjustable parameters of the method. This chart shows that tuning primary parameters will lead to a more desired method and that tuning for secondary tuning parameters leads to even more of a desired result measured with $\mathbf{F}_\beta$ score.

The tuning for $\mathbf{F}_\beta$ score on real data is shown in Figure V.6. This shows a comparison of the methods and how they perform on $\mathbf{F}_\beta$ Score. It also shows the improvement that is made by tuning the primary or all parameters. For the Frequency and Curve Fit methods creeping random search does not improve detection when all parameters are allowed to vary. This is because there was not enough iterations to allow for convergence since the iteration count was the stopping criteria for optimization.

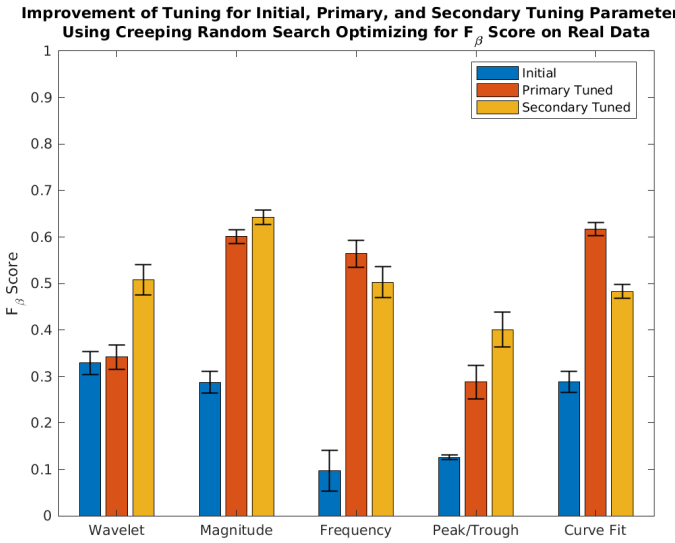

Figure V.6: Tuning results for given methods. The methods were tuned using $\mathbf{F}_\beta$ Score, with $\beta = 0.2$, on real data using Creeping Random Search for optimization. The initial point is given by the standard parameter values for the given method. The primary tuning is the result of tuning for $\mathbf{F}_\beta$ score and only adjusting the primary tuning parameter that would normally be adjusted when a method is used. The secondary tuning is the result of tuning for $\mathbf{F}_\beta$ Score and adjusting all adjustable parameters of the method. This chart shows that tuning primary parameters will lead to a more desired method and that tuning for secondary tuning parameters leads to even more of a desired result measured with $\mathbf{F}_\beta$ score.

The tuning for $\mathbf{F}_\beta$ score on synthetic data is shown in Figure V.7. This shows a comparison of the methods and how they perform on $\mathbf{F}_\beta$ Score. It also shows the improvement that is made by tuning the primary or all parameters. All methods improved when more parameters were tuned.

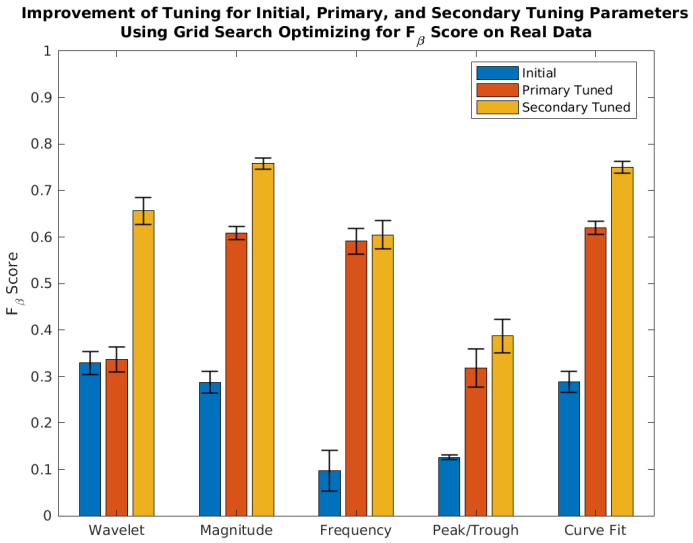

Figure V.7: Tuning results for given methods. The methods were tuned using $\mathbf{F}_\beta$ Score, with $\beta = 0.2$, on real data using univariate grid search for optimization. The initial point is given by the standard parameter values for the given method. The primary tuning is the result of tuning for $\mathbf{F}_\beta$ score and only adjusting the primary tuning parameter that would normally be adjusted when a method is used. The secondary tuning is the result of tuning for $\mathbf{F}_\beta$ Score and adjusting all adjustable parameters of the method. This chart shows that tuning primary parameters will lead to a more desired method and that tuning for secondary tuning parameters leads to even more of a desired result measured with $\mathbf{F}_\beta$ score.

V.4 Conclusion

Using these optimization approaches methods for extracting bursts can be improved. Using the metric of $\mathbf{F}_\beta$ Score each method was shown to improve from the initial conditions. Optimizing for this metric means that the methods with the new tuning parameters will extract more true bursts and fewer false bursts.

This work suggests that to obtain better results for burst detection and characterization methods tuning should be performed over all possible tuning parameters. The metric to use as the objective function should be the $\mathbf{F}_\beta$ score.

Future work in this optimization could include using different data sets for the optimization. This could show how sensitive the tuning is to the data. It may show that tuning is data set specific and that optimization should be done for each different data set.

Future work could also include using different kinds of metrics that would tune for specific requirements such as a good frequency match between detected and ground truth events. In this case a metric that focuses on frequency match such as the Frequency Metric metric might be used. Other metrics could be combined in an averaging scheme so multiple metrics could make a combined metric like CAM.

CHAPTER VI

Conclusion

Detecting and characterizing oscillations in brain signals is an active field of study. Tuning the algorithms used to do so is a non-trivial problem, especially for secondary tuning parameters that are set ad-hoc at design time. This work presents a rigorous study of performance metrics that are useful for measuring how well an algorithm is performing its desired tasks with different performance metrics being important for different chosen applications. A new performance metric useful for tuning algorithms when phase characterization is important is also presented, the CAM.

This work then applies these performance metrics to automated tuning of primary and secondary tuning parameters. In conclusion, the tools presented in this work enable improved performance of oscillation detection and characterization methods and make it practical to tune algorithms that have a large number of tuning parameters. The design of new detection and characterization methods using performance metrics to cluster, accept or reject detected events is identified as an avenue for future work.

This work suggests that to obtain better results for burst detection and characterization methods tuning should be performed over all possible tuning parameters. The metric to use as the objective function should be the $\mathbf{F}_\beta$ score.

BIBLIOGRAPHY

[1] S. Zanos, I. Rembado, D. Chen, and E. E. Fetz, "Phase-locked stimulation during cortical beta oscillations produces bidirectional synaptic plasticity in awake monkeys," *Current Biology*, vol. 28, no. 16, pp. 2515–2526, 2018.

[2] M. Seo, E. Lee, and B. B. Averbeck, "Action selection and action value in frontal-striatal circuits," *Neuron*, vol. 74, no. 5, pp. 947–960, 2012.

[3] T. Schönberg, N. D. Daw, D. Joel, and J. P. O'Doherty, "Reinforcement learning signals in the human striatum distinguish learners from nonlearners during reward-based decision making," *Journal of Neuroscience*, vol. 27, no. 47, pp. 12860–12867, 2007.

[4] M. Bear, B. Connors, and M. A. Paradiso, *Neuroscience: Exploring the brain*. Jones & Bartlett Learning, LLC, 2020.

[5] I. B. Levitan and L. K. Kaczmarek, *The neuron: cell and molecular biology*. Oxford University Press, USA, 2015.

[6] B. Hille, "Ionic channels in excitable membranes. current problems and biophysical approaches," *Biophysical Journal*, vol. 22, no. 2, pp. 283–294, 1978.

[7] A. Destexhe and C. Bedard, "Local field potentials (lfp)," *Encyclopedia of Computational Neuroscience*, 03 2015.

[8] M. A. Lebedev and M. A. Nicolelis, "Brain-machine interfaces: From basic science to neuroprostheses and neurorehabilitation," *Physiological reviews*, vol. 97, no. 2, pp. 767–837, 2017.

[9] E. Niedermeyer, F. L. Da Silva, E. Niedermeyer, and F. H. Lopes Da Silva, *Electroencephalography : Basic Principles, Clinical Applications, and Related Fields*. Philadelphia, UNITED STATES: Wolters Kluwer, 2004.

[10] A. Destexhe, D. Contreras, and M. Steriade, "Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states," *The Journal of neuroscience : the official journal of the Society for Neuroscience*, vol. 19, no. 11, pp. 4595–4608, 1999. 10341257[pmid] PMC6782626[pmcid].

[11] S. Katzner, I. Nauhaus, A. Benucci, V. Bonin, D. L. Ringach, and M. Carandini, "Local origin of field potentials in visual cortex," *Neuron*, vol. 61, no. 1, pp. 35–41, 2009. 19146811[pmid] PMC2730490[pmcid] S0896-6273(08)01006-4[PII].

[12] M. Lundqvist, J. Rose, P. Herman, S. L. Brincat, T. J. Buschman, and E. K. Miller, "Gamma and beta bursts underlie working memory," *Neuron*, vol. 90, no. 1, pp. 152–164, 2016.

[13] P. J. Uhlhaas and W. Singer, "Abnormal neural oscillations and synchrony in schizophrenia," *Nature reviews neuroscience*, vol. 11, no. 2, pp. 100–113, 2010.

[14] T. A. Whitten, A. M. Hughes, C. T. Dickson, and J. B. Caplan, "A better oscillation detection method robustly extracts eeg rhythms across brain state changes: the human alpha rhythm as a test case," *Neuroimage*, vol. 54, no. 2, pp. 860–874, 2011.

[15] M. A. Sherman, S. Lee, R. Law, S. Haegens, C. A. Thorn, M. S. Hämäläinen, C. I. Moore, and S. R. Jones, "Neural mechanisms of transient neocortical beta rhythms: Converging evidence from humans, computational modeling, monkeys, and mice," *Proceedings of the National Academy of Sciences*, vol. 113, no. 33, pp. E4885–E4894, 2016.

[16] J. Feingold, D. J. Gibson, B. DePasquale, and A. M. Graybiel, "Bursts of beta oscillation differentiate postperformance activity in the striatum and motor cortex of monkeys performing movement tasks," *Proceedings of the National Academy of Sciences*, vol. 112, no. 44, pp. 13687–13692, 2015.

[17] H. Cagnan, D. Pedrosa, S. Little, A. Pogosyan, B. Cheeran, T. Aziz, A. Green, J. Fitzgerald, T. Foltynie, and P. Limousin, "Stimulating at the right time: phase-specific deep brain stimulation," *Brain*, vol. 140, no. 1, pp. 132–145, 2017.

[18] B. Picinbono, "On instantaneous amplitude and phase of signals," *IEEE Transactions on signal processing*, vol. 45, no. 3, pp. 552–560, 1997.

[19] A. B. Holt, E. Kormann, A. Gulberti, M. Pötter-Nerger, C. G. McNamara, H. Cagnan, M. K. Baaske, S. Little, J. A. Köppen, and C. Buhmann, "Phase-dependent suppression of beta oscillations in parkinson's disease patients," *Journal of Neuroscience*, vol. 39, no. 6, pp. 1119–1134, 2019.

[20] L. L. Chen, R. Madhavan, B. I. Rapoport, and W. S. Anderson, "Real-time brain oscillation detection and phase-locked stimulation using autoregressive spectral estimation and time-series forward prediction," *IEEE transactions on Biomedical Engineering*, vol. 60, no. 3, pp. 753–762, 2013.

[21] E. Blackwood, M.-c. Lo, and S. A. Widge, "Continuous phase estimation for phase-locked neural stimulation using an autoregressive model for signal prediction," in *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 4736–4739, IEEE, 2018.

[22] J. H. Siegle and M. A. Wilson, "Enhancement of encoding and retrieval functions through theta phase-specific manipulation of hippocampus," *elife*, vol. 3, p. e03061, 2014.

[23] O. Peles, U. Werner-Reiss, H. Bergman, Z. Israel, and E. Vaadia, "Phase-specific microstimulation differentially modulates beta oscillations and affects behavior," *Cell Reports*, vol. 30, no. 8, pp. 2555–2566. e3, 2020.

[24] D. Battaglia, A. Witt, F. Wolf, and T. Geisel, "Dynamic effective connectivity of inter-areal brain circuits," *PLoS computational biology*, vol. 8, no. 3, 2012.

[25] K.-i. Amemori, S. Amemori, D. J. Gibson, and A. M. Graybiel, "Striatal microstimulation induces persistent and repetitive negative decision-making predicted by striatal beta-band oscillation," *Neuron*, vol. 99, no. 4, pp. 829–841, 2018.

[26] S. R. Santacruz, E. L. Rich, J. D. Wallis, and J. M. Carmena, "Caudate microstimulation increases value of specific choices," *Current Biology*, vol. 27, no. 21, pp. 3375–3383, 2017.

[27] D. R. Merrill, M. Bikson, and J. G. Jefferys, "Electrical stimulation of excitable tissue: design of efficacious and safe protocols," *Journal of neuroscience methods*, vol. 141, no. 2, pp. 171–198, 2005.

[28] T. J. Freeborn, "A survey of fractional-order circuit models for biology and biomedicine," *IEEE Journal on emerging and selected topics in circuits and systems*, vol. 3, no. 3, pp. 416–424, 2013.

[29] L. A. Geddes, "Historical evolution of circuit models for the electrode-electrolyte interface," *Annals of biomedical engineering*, vol. 25, no. 1, pp. 1–14, 1997. Geddes, L A eng Historical Article Review Ann Biomed Eng. 1997 Jan-Feb;25(1):1-14. doi: 10.1007/BF02738534.

[30] S. Grimnes and O. G. Martinsen, "Cole electrical impedance model-a critique and an alternative," *IEEE transactions on biomedical engineering*, vol. 52, no. 1, pp. 132–135, 2004.

[31] J. E. B. Randles, "Kinetics of rapid electrode reactions," *Discussions of the faraday society*, vol. 1, pp. 11–19, 1947.

[32] E. M. Moraud, G. Tinkhauser, M. Agrawal, P. Brown, and R. Bogacz, "Predicting beta bursts from local field potentials to improve closed-loop dbs paradigms in parkinson's patients," in *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 3766–3796, IEEE, IEEE, 2018.

[33] S. Cole and B. Voytek, "Cycle-by-cycle analysis of neural oscillations," *Journal of Neurophysiology*, vol. 122, no. 2, pp. 849–861, 2019. PMID: 31268801.

[34] S. Shirinpour, I. Alekseichuk, K. Mantell, and A. Opitz, "Experimental evaluation of methods for real-time eeg phase-specific transcranial magnetic stimulation," *Journal of neural engineering*, vol. 17, no. 4, p. 046002, 2020. Shirinpour, Sina Alekseichuk, Ivan Mantell, Kathleen Opitz, Alexander eng R01 NS109498/NS/NINDS NIH HHS/ Research Support, Non-U.S. Gov't Research Support, N.I.H., Extramural England J Neural Eng. 2020 Jul 13;17(4):046002. doi: 10.1088/1741-2552/ab9dba.

[35] J. B. Caplan, J. R. Madsen, S. Raghavachari, and M. J. Kahana, "Distinct patterns of brain oscillations underlie two basic parameters of human maze learning," *Journal of neurophysiology*, vol. 86, no. 1, pp. 368–380, 2001.

[36] V. Kanta, D. Pare, and D. B. Headley, "Closed-loop control of gamma oscillations in the amygdala demonstrates their role in spatial memory consolidation," *Nature communications*, vol. 10, no. 1, pp. 1–14, 2019.

[37] D. H. Von Seggern, *CRC standard curves and surfaces with mathematica*. CRC Press, 2017.

[38] N. A. M. Norani, W. Mansor, and L. Y. Khuan, "A review of signal processing in brain computer interface system," in *2010 IEEE EMBS Conference on Biomedical Engineering and Sciences (IECBES)*, pp. 443–449, IEEE, IEEE, 2010.

[39] S. L. Schmidt, J. J. Peters, D. A. Turner, and W. M. Grill, "Continuous deep brain stimulation of the subthalamic nucleus may not modulate beta bursts in patients with parkinson's disease," *Brain Stimulation*, vol. 13, no. 2, pp. 433–443, 2020.

[40] F. Mansouri, K. Dunlop, P. Giacobbe, J. Downar, and J. Zariffa, "A fast eeg forecasting algorithm for phase-locked transcranial electrical stimulation of the human brain," *Frontiers in neuroscience*, vol. 11, p. 401, 2017.

[41] A. Bashashati, M. Fatourechi, R. K. Ward, and G. E. Birch, "A survey of signal processing algorithms in brain–computer interfaces based on electrical brain signals," *Journal of Neural engineering*, vol. 4, no. 2, p. R32, 2007.

[42] G. Regalia, S. Coelli, E. Biffi, G. Ferrigno, and A. Pedrocchi, "A framework for the comparative assessment of neuronal spike sorting algorithms towards more accurate off-line and on-line microelectrode arrays data analysis," *Computational intelligence and neuroscience*, vol. 2016, 2016.

[43] C. Thomas and T. Womelsdorf, "Frequency stability detection." 2020.

[44] "Attention Circuits Control Laboratory – wlBurst v2 Library." https://github.com/att-circ-contrl/wlBurst_v2. Accessed: 2020-09-25.

[45] 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," *Journal of neurophysiology*, vol. 116, no. 2, pp. 306–321, 2016.

[46] B. Boashash, *Time-Frequency Methodologies in Neurosciences*, book section 16, pp. 915–966. Academic Press, 2016.

[47] B. Boashash, G. Azemi, and N. A. Khan, "Principles of time–frequency feature extraction for change detection in non-stationary signals: Applications to newborn eeg abnormality detection," *Pattern Recognition*, vol. 48, no. 3, pp. 616–627, 2015.

[48] P. Y. Ktonas and N. Papp, "Instantaneous envelope and phase extraction from real signals: theory, implementation, and an application to eeg analysis," *Signal Processing*, vol. 2, no. 4, pp. 373–385, 1980.

[49] H. Huang and J. Pan, "Speech pitch determination based on hilbert-huang transform," *Signal Processing*, vol. 86, no. 4, pp. 792–803, 2006.

[50] M. Feldman, "Hilbert transform in vibration analysis," *Mechanical systems and signal processing*, vol. 25, no. 3, pp. 735–802, 2011.

[51] G. Tinkhauser, A. Pogosyan, S. Little, M. Beudel, D. M. Herz, H. Tan, and P. Brown, "The modulatory effect of adaptive deep brain stimulation on beta bursts in parkinson's disease," *Brain*, vol. 140, no. 4, pp. 1053–1067, 2017.

[52] A. Teolis and J. J. Benedetto, *Computational signal processing with wavelets*, vol. 182. Springer, 1998.

[53] I. Simonovski and M. Boltežar, "The norms and variances of the gabor, morlet and general harmonic wavelet functions," *Journal of Sound and Vibration*, vol. 264, no. 3, pp. 545–557, 2003.

[54] S. Chandran KS, C. S. Seelamantula, and S. Ray, "Duration analysis using matching pursuit algorithm reveals longer bouts of gamma rhythm," *Journal of neurophysiology*, vol. 119, no. 3, pp. 808–821, 2018.

[55] J. M. Lilly and S. C. Olhede, "Higher-order properties of analytic wavelets," *IEEE Transactions on Signal Processing*, vol. 57, no. 1, pp. 146–160, 2008.

[56] I. Tal and M. Abeles, "Temporal accuracy of human cortico-cortical interactions," *Journal of neurophysiology*, vol. 115, no. 4, pp. 1810–1820, 2016.

[57] B. Boashash, G. Azemi, and J. M. O'Toole, "Time-frequency processing of nonstationary signals: Advanced tfd design to aid diagnosis with highlights from medical applications," *IEEE signal processing magazine*, vol. 30, no. 6, pp. 108–119, 2013.

[58] S. P. Burns, D. Xing, M. J. Shelley, and R. M. Shapley, "Searching for autocoherence in the cortical network with a time-frequency analysis of the local field potential," *Journal of Neuroscience*, vol. 30, no. 11, pp. 4033–4047, 2010.

[59] S. P. Burns, D. Xing, and R. M. Shapley, "Is gamma-band activity in the local field potential of v1 cortex a "clock" or filtered noise?," *Journal of Neuroscience*, vol. 31, no. 26, pp. 9658–9664, 2011.

[60] D. Xing, Y. Shen, S. Burns, C.-I. Yeh, R. Shapley, and W. Li, "Stochastic generation of gamma-band activity in primary visual cortex of awake and anesthetized monkeys," *Journal of Neuroscience*, vol. 32, no. 40, pp. 13873–13880a, 2012.

[61] J. M. Mendel and K. S. Fu, *Adaptive, learning, and pattern recognition systems; theory and applications*, vol. 66 of *Mathematics in Science and Engineering*. New York: Academic Press, 1970.

[62] R. White Jr, "A survey of random methods for parameter optimization," *Simulation*, vol. 17, no. 5, pp. 197–205, 1971.

[63] A. Zhigljavsky and A. Zilinskas, *Stochastic Global Optimization*, vol. 9 of *Springer Optimization and Its Applications*. New York, NY, UNITED STATES: Springer, 1 ed., 2007.

[64] M. Friedman and L. J. Savage, *Planning Experiments Seeking Maxima*, book section 13, pp. 363–372. New York and London: McGraw-Hill, New York, NY, 1947.

[65] S. H. Brooks, "A discussion of random methods for seeking maxima," *Operations Research*, vol. 6, no. 2, pp. 244–251, 1958.

[66] Z. B. Zabinsky, "Stochastic methods for practical global optimization," *Journal of Global Optimization*, vol. 13, no. 4, pp. 433–444, 1998.

[67] C. C. Y. Dorea, "Expected number of steps of a random optimization method," *Journal of Optimization Theory and Applications*, vol. 39, no. 2, pp. 165–171, 1983.

[68] F. J. Solis and R. J.-B. Wets, "Minimization by random search techniques," *Mathematics of operations research*, vol. 6, no. 1, pp. 19–30, 1981.

[69] N. Baba, "Convergence of a random optimization method for constrained optimization problems," *Journal of Optimization Theory and Applications*, vol. 33, no. 4, pp. 451–461, 1981.

[70] D. Andreev, S. Bozhokin, I. Venevtsev, and K. Zhunusov, "Gabor transform and continuous wavelet transform for model pulsed signals," *Technical Physics*, vol. 59, no. 10, pp. 1428–1433, 2014.

[71] S. M. Ross, *Introduction to probability models*. Academic press, 2014.

[72] I. Tal, S. Neymotin, S. Bickel, P. Lakatos, and C. E. Schroeder, "Oscillatory bursting as a mechanism for temporal coupling and information coding," *Frontiers in Computational Neuroscience*, vol. 14, no. 82, 2020.