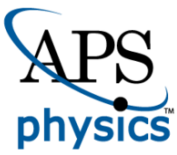

Enabling high confidence detections of gravitational-wave bursts

Tyson B. Littenberg, Jonah B. Kanner, Neil J. Cornish, and Margaret Millhouse

Phys. Rev. D **94**, 044050 — Published 25 August 2016

DOI: 10.1103/PhysRevD.94.044050

Enabling high confidence detections of gravitational-wave bursts

Tyson B. Littenberg
*Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA) & Department of Physics and Astronomy, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208**

Jonah B. Kanner
LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, USA

Neil J. Cornish and Margaret Millhouse
Department of Physics, Montana State University, Bozeman, MT 59717, USA

Extracting astrophysical information from gravitational wave detections is a well-posed problem and thoroughly studied when detailed models for the waveforms are available. However, one motivation for the field of gravitational wave astronomy is the potential for new discoveries. Recognizing and characterizing unanticipated signals requires data analysis techniques which do not depend on theoretical predictions for the gravitational waveform. Past searches for short-duration un-modeled gravitational wave signals have been hampered by transient noise artifacts, or "glitches," in the detectors. We have put forth the `BayesWave` algorithm to differentiate between generic gravitational wave transients and glitches, and to provide robust waveform reconstruction and characterization of the astrophysical signals. Here we study `BayesWave`'s capabilities for rejecting glitches while assigning high confidence to detection candidates through analytic approximations to the Bayesian evidence. Analytic results are tested with numerical experiments by adding simulated gravitational wave transient signals to LIGO data collected between 2009 and 2010 and found to be in good agreement.

I. INTRODUCTION

Among the most intriguing promises of the nascent field of gravitational wave (GW) astronomy is the discovery of never-before anticipated sources. As the LIGO [1] and Virgo [2] observatories continue to improve their detection sensitivity, so too must the methods used to search through the data looking for the unexpected. Making a claim of a significant discovery requires exceptional evidence. In the field of particle physics, a common practice for declaring detection of a new particle is a "5-sigma" level of confidence, meaning that there is probability of less than 3×10^{-7} of the observation arising from sources other than the claimed discovery.

Having detailed theoretical predictions for the gravitational wave signal helps reduce the false alarm (or false positive) rate due [3–6] but searches for generic signals, known as GW bursts, have to confront non-Gaussian noise artifacts, or "glitches," in order to indentify astrophysical signals at high confidence (e.g. Ref. [7]). Background distributions for burst searches, determined by time-shifting the data from multiple detectors so that no coherent gravitational signals are in the data, show a long tail to high signal to noise ratio (SNR), meaning that even a very strong gravitational wave signal would be consistent with having arisen from a glitch.

In preparation for the advanced detector era several new approaches to the burst detection problem have been developed. Thrane and Coughlin [10] have demonstrated the capability to make high-confidence detections of long-duration ($\mathcal{O}(10)s$) burst signals in non-stationary, non-Gaussian noise by searching for excess power found along parameterized curves through a time-frequency representation of the data. In an independent effort, the Bayesian parameter estimation analysis library `LALInference` [11], originally designed for the characterization of compact binary signals, has been adapted for burst analyses by using a sine-Gaussian waveform as the gravitational wave template [12, 13]. `LALInference` differentiates between signals and glitches using a "coherence test" where the "coherent" signal hypothesis uses a template-based analysis assuming the data streams from multiple detectors contain a coherent gravitational wave signal while the "incoherent" glitch hypothesis treats each data stream independently. The incoherent model uses the same template waveform as the signal model but optimizes its parameters independently for each detectors' data [14].

Recently we proposed `BayesWave`–a Bayesian algorithm to follow-up short duration ($\lesssim$ 1 s) candidate gravitational wave transient events, separate signals from glitches, and provide robust signal characterization for arbitrary burst waveforms [15]. `BayesWave` uses a variable-dimension model for signals and/or glitches enabling the analyses to adapt the complexity of the waveform model to match what is present in the data instead of imposing a template waveform and searching for best fit parameters. For a detection candidate `BayesWave` computes the relative evidence of the event being produced by a GW signal, an instrument artifact, or statistical fluctuations of the detector's Gaussian noise. In the event that the candidate is of astrophysical origin,

* tyson.littenberg@ligo.org

BayesWave also produces posterior distributions for the source sky-location and orientation, accurate waveform reconstruction, and metrics to characterize the signal such as duration, bandwidth, signal energy, etc. In all instances, BayesWave characterizes the instrument behavior including spectral estimation for the background Gaussian noise and glitch reconstructions which can then be used to feedback into the never-ending effort to improve the interferometers' performance. Analysis of the Gaussian component of the instrument noise is handled by BayesWave's sibling algorithm, BayesLine [16]. During the first Advanced LIGO observing run BayesWave is being utilized as a follow-up analysis to candidate and background events found by the coherent WaveBurst algorithm [17].

In this paper we will demonstrate BayesWave's potential by analyzing data from the sixth LIGO science run (S6) which took place from 2009-2010. Our results are achieved by analyzing data known to contain glitches which contributed to the long-tailed background distribution for the burst search, and by adding simulated gravitational wave signals to detector noise. In addition to this study using archived data, we present an analytical framework for understanding the performance of the pipeline. A companion paper uses BayesWave and the flagship burst search algorithm, coherent WaveBurst [17], in an end-to-end demonstration of how burst detection efficiency is improved by the joint analysis [18].

In section II we briefly describe the BayesWave algorithm, Bayesian model selection, and our model for the data. In section III we go through a simple analytic calculation to give insight into how BayesWave is able to distinguish signals and glitches, and use BayesWave's performance on simulated signals added to real data to support the analytic approximations. Section IV uses the intuition built from the analytics to estimate background rates for glitches to be considered signals by BayesWave, and connects the Bayes factor to false alarm rates for detections. We summarize the work in section V. The appendix contains a more detailed derivation of the analytic approximation to the evidence.

II. METHOD

Searches for Burst signals have been based on frameworks that employ detection statistics to measure the likelihood that Gaussian noise could produce the data [17, 19–22]. While stationary Gaussian noise is often a good description for LIGO/Virgo data, the approximation breaks down with much higher regularity than the arrival of detectable gravitational waves. Any data analysis method must account for the possibility of non-stationary non-Gaussian noise. Most existing analysis strategies apply various selection cuts to separate glitches from astrophysical signals which are tuned by adjusting thresholds to minimize the estimated background rate of transient noise glitches [7–9].

Bayesian hypothesis testing has been used in searches for GWs from a timing glitch in the Vela pulsar [23] using a damped sinusoid that abruptly starts at times associated with the pulsar timing glitch as the signal model. A recently developed search pipeline [13] uses excess power to identify interesting data segments and a matched-filtering follow-up with a sine-Gaussian template for signal characterization [12]. The "coherent vs. incoherent" Bayes factor is used to distinguish between noise and signal [14].

BayesWave employs a different approach by using a parameterized model for the LIGO/Virgo data, noise and signal included, and forward modeling, i.e. predicting, the detector output. The data model has three distinct components: A gravitational wave signal h that is elliptically polarized and is coherent across the network of detectors; glitches that are independent in each interferometer; and stationary Gaussian noise which is fully characterized by its power spectral density $S_n(f)$ as modeled by BayesLine [16]. At its core, BayesWave is a Markov chain Monte Carlo (MCMC) algorithm [24]. BayesWave uses parallel tempering [25] and thermodynamic integration [26] to compute the evidence for each model. The MCMC implementation and evidence calculation is described in detail in Refs. [15, 16]. For results in this work we utilize an adaptive temperature scheme as suggested in [27].

Because we do not know *a priori* the functional form of glitch or GW burst waveforms, our model for both must be flexible. We use a linear combination of Morlet-Gabor wavelets as our waveform model where the number of wavelets included in the linear combination, N, is itself a model parameter. Each basis function (wavelet) is described by parameter vector $\boldsymbol{\lambda} \rightarrow \{f_0, t_0, A, Q, \varphi_0\}$ with components for central frequency f_0 ; central time t_0 ; amplitude A; quality factor Q; and phase offset φ_0. A wavelet is expressed in the time domain as:

$$\Psi(t; A, f_0, Q, t_0, \phi_0) = Ae^{-\Delta t^2/\tau^2} \cos(2\pi f_0 \Delta t + \varphi_0) \quad (1)$$

where $\tau = Q/(2\pi f_0)$ and $\Delta t = t - t_0$. BayesWave uses a reversible jump Markov chain Monte Carlo [28] to marginalize over the number of wavelets needed for the model to be consistent with the data.

A. Bayesian hypothesis testing or model selection

This work will rely on Bayes' rule for conditional probabilities to make inferences about the measured data d. The conditional probabilities $p(A|B, C)$ are to be understood as the probability density of A given B and C. Quantities with the data d appearing to the left of the vertical bar $p(d|B, C)$ are likelihoods of observing data d given B and C. When the data d appears to the right of the bar $p(A|d, C)$ the quantity is a posterior probability for measuring A given the data and C. If the data d do not appear in $p(A|B)$ the quantity is prior probability density of A given B.

The likelihood that hypothesis $\mathcal{H}$, parameterized by $\boldsymbol{\theta}$, would have produced the data d is calculated by

$$p(d|\mathcal{H}) = \int p(d|\boldsymbol{\theta},\mathcal{H})p(\boldsymbol{\theta}|\mathcal{H})d\boldsymbol{\theta}. \quad (2)$$

Directly integrating Eq 2 is seldom practical and a wide variety of alternative means for arriving at $p(d|\mathcal{H})$ have been devised. Our method of choice for computing the integral in Eq. 2 is thermodynamic integration [26].

Once the evidence has been computed it provides a relative measure of how well one hypothesis is supported by the data over another through the "odds ratio"

$$\mathcal{O}_{0,1} \equiv \frac{p(\mathcal{H}_0)}{p(\mathcal{H}_1)}\frac{p(d|\mathcal{H}_0)}{p(d|\mathcal{H}_1)} = \frac{p(\mathcal{H}_0)}{p(\mathcal{H}_1)}\mathcal{B}_{0,1} \quad (3)$$

where $p(\mathcal{H})$ is the prior probability for the hypothesis and $\mathcal{B}_{0,1}$ is the likelihood ratio, or "Bayes factor," for the two hypotheses.

In what follows, we will explicitly write out the model evidence, and ultimately the Bayes factor, for comparing the signal and glitch hypotheses. Throughout this derivation we will use $\boldsymbol{\theta}$ as the vector of all model parameters. For the glitch hypothesis, $\boldsymbol{\theta}$ is separated into independent sets of *intrinsic parameters* $\boldsymbol{\lambda}_i^I$ which encode the shape of the sine-Gaussian waveform. Superscripts on $\boldsymbol{\theta}$ identify which detectors data the wavelet is applied while subscripts identify the different wavelets in the linear combination. For the signal hypothesis, there is a common set of intrinsic parameters $\boldsymbol{\lambda}^\oplus$ applied at the earth geocenter and projected on to the network, where again the subscripts indicate which wavelet in the series. Encoding the projection are the *extrinsic parameters* $\boldsymbol{\Omega}$.

B. Modeling signals versus glitches

Consider a GW network consisting of the two LIGO detectors in Hanford, WA (H) and Livingston, LA (L). For a candidate event `BayesWave` calculates the Bayesian evidence for each of three models: signal, glitch, or Gaussian noise. We can then use the Bayes factor between any two models to quantify the degree of supporting evidence for one model over the other. Within each model the likelihood is computed by

$$p(d|\boldsymbol{\theta},\mathcal{H}) \propto \prod_I^{H,L} e^{-\frac{1}{2}(r^I(\boldsymbol{\theta})|r^I(\boldsymbol{\theta}))} \quad (4)$$

where r^I is the residual of the data in detector I minus the signal or glitch model, $(a|b) \equiv \frac{2}{T}\int \frac{\tilde{a}^*(f)\tilde{b}(f)+\tilde{a}(f)\tilde{b}^*(f)}{S_n(f)}df$ is the noise weighted inner product, $S_n(f)$ is the noise power spectral density estimated from the data by `BayesLine`, and T is the duration of the data. This work will focus only on examples where we need to distinguish between the signal and glitch models. We assume either will be preferred over the Gaussian noise model.

1. $\mathcal{H}_0$: The glitch model ($\mathcal{G}$):

The data $d^I = n^I + g^I$ contain Gaussian noise n and glitches g independent in each detector I. The parameters $\boldsymbol{\theta} \to [\boldsymbol{\lambda}^H \cup \boldsymbol{\lambda}^L]$ are comprised of independent sets of intrinsic parameters

$$\boldsymbol{\lambda}^I \to [\boldsymbol{\lambda}_0 \cup \boldsymbol{\lambda}_1 \cup \cdots \cup \boldsymbol{\lambda}_{N^I}]$$

which determine the shape of each wavelet. The glitch model is computed for each detector as an independent linear combination of wavelets

$$g(\boldsymbol{\lambda}^I, N^I) = \sum_i^{N^I} \tilde{\Psi}(f;\boldsymbol{\lambda}_i^I)$$

where $\tilde{\Psi}(f)$ is the Fourier transform of $\Psi(t)$, N^I is the number of wavelets used in the sum and can take on any value between $[0, N_{\rm max}]$ with the caveat that *at least one wavelet* must be used in the model for the whole network. $N_{\rm max}$ is typically 20. The glitch-model likelihood is computed using Eq. 4 with the residual $r^I(\boldsymbol{\theta}) = d^I - g(\boldsymbol{\lambda}^I, N^I)$

2. $\mathcal{H}_1$: The signal model ($\mathcal{S}$):

The data $d^I = n^I + h^I$ contain Gaussian noise n and an elliptically polarized gravitational wave signal h coherent across the network of detectors. The parameters $\boldsymbol{\theta} \to [\boldsymbol{\lambda}^\oplus \cup \boldsymbol{\Omega}]$ are a common set of intrinsic parameters

$$\boldsymbol{\lambda}^\oplus \to [\boldsymbol{\lambda}_0 \cup \boldsymbol{\lambda}_1 \cup \cdots \cup \boldsymbol{\lambda}_{N^\oplus}]$$

referenced at the center of the Earth and four "extrinsic parameters"

$$\boldsymbol{\Omega} \to [\theta, \phi, \psi, \epsilon]$$

which define the sky-location θ, ϕ, the polarization angle ψ and an ellipticity parameter ϵ relating the two gravitational wave polarizations h_+ and $h_\times$. The signal-model likelihood is computed using Eq. 4 with the residual

$$r^I(\boldsymbol{\theta}) = d^I - h^I(f;\boldsymbol{\lambda}^\oplus, N^\oplus, \boldsymbol{\Omega})$$

The geocenter signal wavelets are projected onto the network using each detector's unique time delay operators $\Delta t(\theta,\phi)$, and antenna beam pattern response functions $F^+(\theta,\phi,\psi)$, $F^\times(\theta,\phi,\psi)$ [29]:

$$\begin{aligned} h^I(f;\boldsymbol{\lambda}^\oplus, N^\oplus, \boldsymbol{\Omega}) &= \left(F_I^+ h_+(f) + F_I^\times h_\times(f)\right) e^{2\pi i f \Delta t_I} \\ h_+(f) &= \sum_i^{N^\oplus} \tilde{\Psi}(f;\boldsymbol{\lambda}_i^\oplus) \\ h_\times(f) &= \epsilon\, h_+(f) e^{i\pi/2}. \end{aligned} \quad (5)$$

III. DISTINGUISHING SIGNALS FROM GLITCHES

While BayesWave uses a computationally expensive numerical integration to compute the evidence for each model, we will build intuition for how BayesWave successfully distinguishes signals from glitches using the Laplace approximation to the evidence and several simplifying assumptions about the model and the data. As our results will show, the simple analytic treatment derived here leads to useful approximations for when signals and glitches are distinguishable and in forecasting the most significant background event. A more detailed derivation and discussion of the Laplace approximation to BayesWave's signal and glitch model evidence can be found in the appendix.

A. Laplace-Fisher approximation to the evidence

If an event has enough SNR to be a strong candidate for detection (SNR $\equiv \sqrt{(h|h)} \gtrsim 10$) the integrand of Equation (2) will be sharply peaked around the maximum a posteriori (MAP) parameter values of the model $\boldsymbol{\theta}_{\rm MAP}$. The evidence can be estimated as

$$p(d|\mathcal{H}) \simeq p(d|\boldsymbol{\theta}_{\rm MAP},\mathcal{H})p(\boldsymbol{\theta}_{\rm MAP}|\mathcal{H})(2\pi)^{D/2}\sqrt{\det C_{\boldsymbol{\theta}}} \quad (6)$$

which is the product of the MAP likelihood $p(d|\boldsymbol{\theta}_{\rm MAP},\mathcal{H})$, the prior $p(\boldsymbol{\theta}_{\rm MAP}|\mathcal{H})$ evaluated at the MAP parameters, and the determinant of the full parameter covariance matrix $C_{\boldsymbol{\theta}}$ which is a measure of the posterior volume. D is the dimension of the model. The covariance matrix $C_{\boldsymbol{\theta}}$ can be approximated by the inverse of the Fisher information matrix $\Gamma_{\boldsymbol{\theta}}$, and we replace $\sqrt{\det C_{\boldsymbol{\theta}}}$ with $1/\sqrt{\det \Gamma_{\boldsymbol{\theta}}}$.

The $p(\boldsymbol{\theta}_{\rm MAP}|\mathcal{H})(2\pi)^{D/2}\sqrt{\det C_{\boldsymbol{\theta}}}$ term is the "Occam factor" that penalizes the likelihood by the model's size. If two models achieve the same likelihood the Occam factor, and therefore the evidence, will be smaller for the model that requires more (constrained) parameters to achieve that fit. Consider a simple model with a single parameter x and uniform prior over an interval V_x. The covariance matrix is simply the variance of the likelihood σ_x^2. In this case the Occam factor is proportional to σ_x/V_x which leads to a simple, intuitive, interpretation: *The Occam factor is the fraction of the prior taken up by the posterior.* We will return to this interpretation when predicting the most significant background event for BayesWave.

For the glitch or signal model, the expectation value for the intrinsic parameter log likelihood is proportional to [30]

$$\ln p(\boldsymbol{\lambda}_{\rm MAP}|\mathcal{H}) \sim \frac{\rm SNR^2}{2} + \frac{D}{2}. \quad (7)$$

For uniform priors $p(\boldsymbol{\lambda}_{\rm MAP}|\mathcal{H}) = 1/V_{\boldsymbol{\lambda}}$ where $V_{\boldsymbol{\lambda}}$ is the volume of the intrinsic parameter space. BayesWave uses uniform priors for all but the amplitude parameter, with $p(A)$ a function of the wavelet's SNR [15]. For simplicity we will neglect the parameter-dependence of the amplitude prior in favor of the simpler $1/V$ scaling. A similar but more detailed derivation including the true amplitude used by BayesWave can be found in the appendix.

The determinant of the intrinsic parameter Fisher matrix for a single wavelet is

$$\det \Gamma_{\boldsymbol{\lambda}} = \frac{\pi^2 \rm SNR^{10}}{2Q^2}. \quad (8)$$

If we assume little overlap between wavelets in the parameter space the correlations between wavelet parameters are negligible and the Fisher matrix is block diagonal. The determinant for the full covariance matrix with N wavelets is then [15]

$$\sqrt{\det C_{\boldsymbol{\lambda}}} \approx \prod_n^N \frac{\sqrt{2}Q_n}{\pi \rm SNR_n^5}. \quad (9)$$

Neglecting the extrinsic parameters for the signal model, and the BayesLine parameters which are common to all models, the dimension $D = 5N$ where N is the number of wavelets used in the fit. To simplify the expression we define $\bar{Q}_n \equiv (2\pi)^{5/2}\frac{\sqrt{2}Q_n}{\pi}$ to absorb the $(2\pi)^{D/2}$ and additional factors of 2 and π. Now the log evidence becomes

$$\log p(d|\mathcal{H}) \simeq \frac{\rm SNR^2}{2} + \frac{5N}{2} - N\log(V_{\boldsymbol{\lambda}}) + \sum_n^N \frac{\bar{Q}_n}{\rm SNR_n^5}. \quad (10)$$

From this expression we see that the Bayes factor for either the glitch or signal model versus the Gaussian noise model goes as $\mathcal{O}\left(\rm SNR^2\right)$.

For the glitch model, the prior and posterior volume terms are summed over the number of detectors (IFO) in the network. The signal model, on the other hand, picks up an additional $D_{\boldsymbol{\Omega}}/2$ and Occam factor term $\log\sqrt{\det C_{\boldsymbol{\Omega}}}/V_{\boldsymbol{\Omega}}$ for the extrinsic parameters which govern the projection of the signal onto the network. $D_{\boldsymbol{\Omega}}$ is the extrinsic parameter dimension, $C_{\boldsymbol{\Omega}}$ is the signal parameter covariance matrix, and $V_{\boldsymbol{\Omega}}$ is the extrinsic parameter prior volume. Including these details into Eq. 10 we find the log evidence for the glitch and signal models is

$$\log p(d|\mathcal{G}) \simeq \frac{\mathrm{SNR}^2}{2} + \sum_i^{\mathrm{IFO}} \frac{5N^i}{2} - \sum_i^{\mathrm{IFO}} N^i \log(V_{\boldsymbol{\lambda}}) + \sum_i^{\mathrm{IFO}} \sum_n^{N^i} \frac{\bar{Q}_n^i}{(\mathrm{SNR}_n^i)^5}$$

$$\log p(d|\mathcal{S}) \simeq \frac{\mathrm{SNR}^2}{2} + \frac{5N}{2} - N \log(V_{\boldsymbol{\lambda}}) + \sum_n^N \frac{\bar{Q}_n}{\mathrm{SNR}_n^5} + \left[\frac{D_{\boldsymbol{\Omega}}}{2} + \log \frac{\sqrt{\det C_{\boldsymbol{\Omega}}}}{V_{\boldsymbol{\Omega}}}\right] \tag{11}$$

respectively, where $\mathrm{SNR}^2 = \sum_i^{\mathrm{IFO}} \mathrm{SNR}_i^2$ and the extrinsic parameter dimension $D_{\boldsymbol{\Omega}} = 4$ while the prior volume for extrinsic parameters is $4\pi^2$.

B. Two detector network

Consider a fairly loud gravitational wave signal in the two detector LIGO network. The optimal extrinsic parameters for detection will result in similar signal strength in each of the interferometers such that $\mathrm{SNR}_n^I \approx \mathrm{SNR}_n/\sqrt{2}$ where the index n is for each wavelet and the index I is for each detector. For such events the different hypotheses will use similar wavelets, so each Q_n^I for the glitch model will correspond to a Q_n for the signal mode, but because it treats each detector independently, the glitch model will need two copies of the linear combination–one for each interferometer. One final simplifying assumption is that the signal to noise ratio of each wavelet is the same: $\mathrm{SNR}_n \approx \overline{\mathrm{SNR}} = \mathrm{SNR}/\sqrt{N}$.

Substituting these simplifications into Eq. 11 we arrive at a simple expression for the log Bayes factor $\log \mathcal{B}_{\mathcal{S},\mathcal{G}} = \log p(d|\mathcal{S}) - \log p(d|\mathcal{G})$:

$$\log \mathcal{B}_{\mathcal{S},\mathcal{G}} \simeq \frac{5N}{2} + N \log V_{\boldsymbol{\lambda}} + 5N \log(\overline{\mathrm{SNR}}) - \sum_n^N \log \bar{Q}_n + \left[2 + \log \frac{\sqrt{\det C_{\boldsymbol{\Omega}}}}{4\pi^2}\right] \tag{12}$$

and immediately see that $\log \mathcal{B}_{\mathcal{S},\mathcal{G}} \sim \mathcal{O}(N \log \mathrm{SNR})$. As a consequence, at fixed SNR, waveform morphologies that require more wavelets to reconstruct have a higher likelihood of being classified as signals. This is an important departure from traditional SNR-based ranking statistics. The Bayes factor computed by `BayesWave` is more sensitive to signal complexity than signal strength. Heuristically, the $\log \mathcal{B}_{\mathcal{S},\mathcal{G}}$ naturally encodes how increasingly unlikely it is for the detectors to simultaneously and coherently produce glitches with non-trivial time-frequency structure. This is a significant difference from existing Burst pipelines which put greater emphasis on signal strength in forming their detection statistic, and are thus hamstrung by the detectors' tendency to produce loud noise transients at a higher rate than the universe supplies gravitational wave signals. We find this fundamental difference allows `BayesWave` to assign detection candidates high confidence in data prone to loud glitches while existing pipelines have not.

C. Single wavelet examples in simulated noise

To verify the predictions from the Laplace approximation we used `BayesWave` to recover simulated sine Gaussian gravitational wave signals in Gaussian noise, drawing waveform parameters from our prior distributions: $f \in [16, 512]$ Hz, $t \in [-0.5, 0.5]$ s, $Q \in [3, 40]$, $\varphi \in [0, 2\pi]$ rad, $\cos\theta \in [-1, 1]$, $\phi \in [0, 2\pi]$ rad, $\psi \in [0, \pi/2]$ rad, $\epsilon \in [-1, 1]$, and amplitude drawn from the distribution described in the appendix and Ref. [15]. For this study we analyze segments of LIGO data collected during the sixth science run which took place from 2009-2010 in which we have purposefully added GW signals. The priors used for this analysis reflect what is being used for low-frequency triggers in the first advanced LIGO observing run (O1) during which `BayesWave` relies on the coherent WaveBurst pipeline to provide the segments of data which warrant follow-up analysis (for details see Refs. [17, 18]).

`BayesWave` calculates Bayes factors for each combination of models along with an estimate of the error in that calculation, using thermodynamic integration. We do not anticipate the agreement between numerical simulations and the analytic approximations to be perfect. Many of the approximations we have made along the way to arrive at Eqs. 10 and 12 are known to be inadequate for the gravitational wave detection problem [31], particularly our use of the covariance matrix to estimate the posterior volume and, even more egregiously, our use of the Fisher matrix as the inverse covariance matrix [32].

Fisher matrix approximations are particularly bad for $\det C_{\boldsymbol{\Omega}}$. The extrinsic parameter space is rife with degeneracies between parameters and non-Gaussian, multimodal likelihood distributions which often span the full extent of the prior range. Fisher matrix arguments would predict a $\mathrm{SNR}^{-\mathrm{D}_{\boldsymbol{\Omega}}}$ scaling for the determinant of $C_{\boldsymbol{\Omega}}$ which is much too strong for any burst source in a two detector network. Using numerical experiments to get

a rough understanding of the extrinsic parameter posterior volume, we find an $\mathrm{SNR}^{-\gamma}$ with γ ranging from ~ 1 at low SNRs to ~ 2.5 at SNR ~ 100 (See Fig. 6 in the appendix).

In Figure 1 the left panel shows the glitch to noise (red) and signal to glitch (blue) log Bayes factors as a function of the simulated signals' SNR along with the Laplace approximation predictions. The predicted scaling laws for $N = 1$ signals $\log \mathcal{B}_{[\mathcal{S}/\mathcal{G}],\mathcal{N}} \sim \mathcal{O}\left(\mathrm{SNR}^2\right)$ and $\log \mathcal{B}_{\mathcal{S},\mathcal{G}} \sim \mathcal{O}\left(\log \mathrm{SNR}\right)$ are generally obeyed by the numerical results. The observed agreement reinforces the intuition developed from considering the analytic expressions, and we can be confident that the numerical integration is performing well. The right panel demonstrates `BayesWave`'s glitch rejection capabilities by comparing $\log \mathcal{B}_{\mathcal{S},\mathcal{G}}$ for simulated sine-Gaussian glitches (gray crosses) and signals (blue circles). The glitches were simulated by adding sine-Gaussians to each detector with parameters drawn independently from the prior. Negative $\log \mathcal{B}_{\mathcal{S},\mathcal{G}}$ corresponds to data with higher likelihood for the glitch model.

D. Multiple wavelet examples in real noise

Equation 12 predicts that the Bayes factor grows with SNR more rapidly for waveforms that have more time-frequency structure, thus requiring more wavelets to account for all of the excess power in the data. For astrophysical signals the number of wavelets necessary will not be known a priori, and furthermore will not be constant, depending on the SNR. As the signal strength increases, more detailed structure in the waveforms will be detectable, and more wavelets will be favored by the model selection. Through numerical experiments we find simple relationships for the number of basis functions and the average SNR per wavelet in terms of the true SNR:

$$\begin{aligned} N &\sim 1 + \beta \mathrm{SNR} \\ \overline{SNR} &\sim \alpha \mathrm{SNR}^a \end{aligned} \tag{13}$$

where the coefficients β, α and index a are different for different kinds of signals with $\{\beta{\sim}0, \alpha{\sim}1, a{\sim}1\}$ corresponding to sine Gaussian waveforms and β and α increasing while a decreases with increasing signal complexity (see Fig. 5 in the appendix).

To demonstrate this important feature of `BayesWave` we add simulated gravitational wave signals from different waveform families into real detector data. Figure 2 shows $\log \mathcal{B}_{\mathcal{S},\mathcal{G}}$ as a function of SNR for the different simulations. Red points are sine Gaussian waveforms, blue points correspond to signals from the merger of two 50 $\mathrm{M}_\odot$ black holes modeled using non-spinning Effective One Body (EOB) waveforms [33], and the black points are results from "white noise bursts"–unpolarized, band-limited, white noise signals used to test LIGO/Virgo burst detection pipelines. We can empirically determine that β is larger for more complicated signal morphologies. Results agree well enough with the analytic predictions that the insight gained in the analytic study is applicable, but the Laplace approximation is clearly no substitute for the numerical integration. The large scatter in Bayes factors is due to failings in the Laplace approximation, signals that violate our assumption about roughly equal SNR in each detector, and segments of data that contain both signals and glitches.

It is important to note that the high degree of scatter in the white noise burst results is also to be expected because these signals are unpolarized, while `BayesWave` assumes h_+ and $h_\times$ are related by Eq. 5. In a two detector network we generally cannot reliably measure the GW polarizations independently. Introducing the additional degrees of freedom to independently solve for h_+ and $h_\times$ will hinder our ability to reject glitches because the number of signal model and glitch model parameters will be comparable for a wider variety of waveform morphologies. While there is no reason to expect a priori that GW bursts will be elliptically polarized, selection effects by the detection pipelines which identify segments of data for `BayesWave` to follow-up in a real analysis, and the similar orientation of the LIGO detectors, favor signals which are well approximated by a single polarization (causing many of the degeneracies between extrinsic parameters discussed in the previous section). This assumption will need to be relaxed when more detectors are added to the network, and in future studies we will investigate strategies for optimizing `BayesWave`'s performance on unpolarized detection candidates even in the two detector case.

IV. BACKGROUND ESTIMATION

We have shown that `BayesWave` predictably favors the signal model over the glitch model for simulated GW events, i.e. `BayesWave` is robust against false dismissal of gravitational wave signals. This is only half of the battle: Any useful data analysis procedure must also be robust against false alarms, i.e. misidentifying noise events as being astrophysical signals, and be able to assign significance to a detection. While the right panel of Figure 1 demonstrates how `BayesWave` can reject glitches in the trivial case of random sine-Gaussian waveforms, how it will fair against real glitches, and how to assign significance to candidate events, requires more careful attention.

To understand `BayesWave`'s glitch rejection capabilities, imagine that a glitch waveform in LIGO Hanford (H) is well represented by a linear combination of wavelets with parameters $\boldsymbol{\lambda}^H$ and a coincident (i.e. within the light travel time between detectors) collection of wavelets is found in LIGO Livingston (L). If the signal is astrophysical in nature, the waveform in L must have parameters $\boldsymbol{\lambda}^L$ that are consistent with $\boldsymbol{\lambda}^H$, within the measurement uncertainties $\sqrt{\det C_{\boldsymbol{\lambda}^L}}$ up to the appropriate time, phase, and amplitude shifts due to the geometry of the detector locations and orientations. On

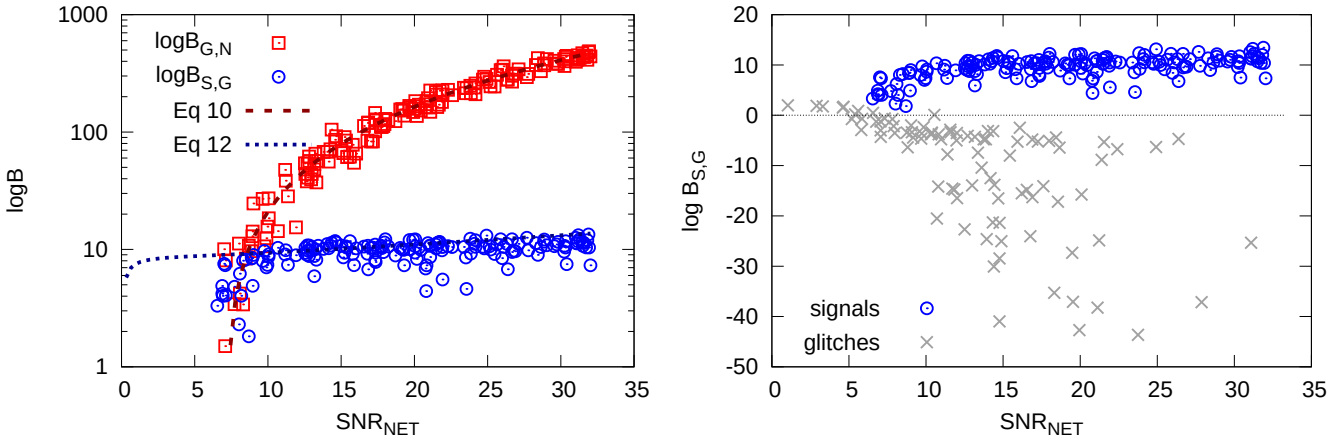

FIG. 1. Left panel: Comparing the numerical results for $\log \mathcal{B}_{\mathcal{G},\mathcal{N}}$ (red squares) and $\log \mathcal{B}_{\mathcal{S},\mathcal{G}}$ (blue circles) to the relevant analytic predictions in Eqs. 10 and 12 (dark red dashed and dark blue dotted lines, respectively) showing good agreement at high SNRs where the Laplace approximation is more valid. Right panel: The $\log \mathcal{B}_{\mathcal{S},\mathcal{G}}$ results from the left panel (blue circles) with $\log \mathcal{B}_{\mathcal{S},\mathcal{G}}$ from a set of glitches simulated by adding independent sine-Gaussian waveforms to each detector (gray crosses). The log Bayes factor shows a clear separation between the simulated signals and glitches. Cases with $\log \mathcal{B}_{\mathcal{S},\mathcal{G}} < 0$ correspond to the glitch model having a higher likelihood.

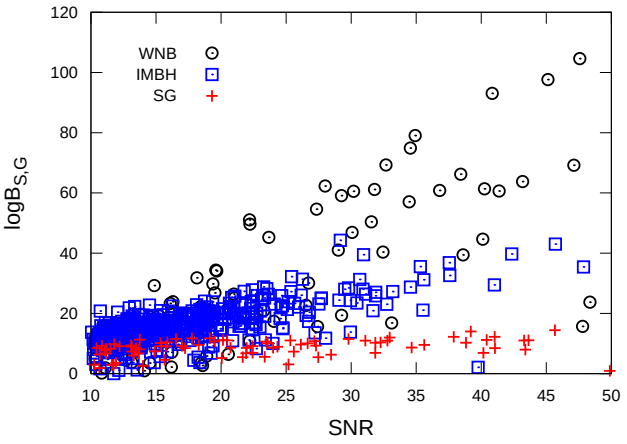

FIG. 2. Bayes factors between the signal and glitch model grow more rapidly with SNR for more complicated waveform morphologies. Sine Gaussian injections (red) only require a single wavelet and thus never achieve particularly high $\log \mathcal{B}_{\mathcal{S},\mathcal{G}}$. Intermediate mass black hole (IMBH) mergers need $N \lesssim 10$ resulting in stronger separation between models (blue), while the white noise burst waveforms (black) have rich time frequency structure, often saturating the prior on N and show the strongest SNR dependence (highest β).

the other hand, if the data represent coincident glitches, then *a priori* there is no reason for the glitch in Livingston to match the parameters in Hanford. Instead, the wavelet in Livingston is chosen at random. One can consider glitches to be random draws from $\boldsymbol{\lambda}$ space and false alarms (glitches that appear as signals) are draws that overlap within the size $\sqrt{\det C_{\boldsymbol{\lambda}^L}}$. If the posteriors do not overlap the data is not consistent with the signal model, i.e. the signal model likelihood will be lower than the glitch model likelihood, and the Bayes Factor will favor the glitch model (c.f. Figure 1).

We can use the same logic to estimate the background rate of glitches that are consistent with the signal model. Assume, in a given two detector data set, there are $\mathcal{N}_{\rm gl}$ coincident glitches. If we assume our signal/glitch model can achieve a perfect match to glitches in the data, the recovered SNR of the signal and the glitch model will be equal when the glitches overlap in the $\boldsymbol{\lambda}$ space and the Bayes factor will be again well approximated by Eq. 12.

Recall the Occam factor is interpreted as the fraction of the prior covered by the posterior $\sim \sqrt{\det C}/V$, i.e. the Occam factor is the size of the "target" the second glitch must hit to be misidentified as a signal. Put another way, a glitch has probability $\sim \sqrt{\det C}/V$ to be consistent with the signal model. Therefore finding a background event with a Bayes factor consistent with Eq 12 will require analysis of something like $(V_{\boldsymbol{\lambda}})^N / \sqrt{\det C_{\boldsymbol{\lambda}}}$ coincident glitches. In our application the Occam factor thus takes on an additional interpretation as the expected number of trials (coincident noise transients) needed for two random glitches to have sufficient overlap in parameter space to look like a signal.

We can loosely turn this into an argument for the maximum Bayes factor–the one that occurs *only once* in a span of LIGO data–as having an Occam factor of $N_{\rm gl}$, i.e. the maximum Bayes factor for a background noise event is

$$\langle \mathcal{B}_{\mathcal{S},\mathcal{G}} \rangle_{\rm background} \lesssim N_{\rm gl}. \tag{14}$$

This limit is not robust. The loudest noise event is obviously in the extreme tail of the background distribution and will therefore fluctuate wildly for different realizations of the data. Nor is this a statement about the population of glitches beyond the assumption that the parameters $\boldsymbol{\lambda}$ are chosen at random for glitches in each detector. It is also important to point out this is may be a conservative estimate. Most glitches are at low SNR in any realistic glitch population, and so low values of the Occam Factor will likely be much more common than high values.

We use our estimate of the most significant background event to approximate the false alarm rate. To do so we need to know the rate of coincident glitches, $R_{\rm gl}$, which is a carefully studied quantity within LIGO. The single detector glitch rate was known during S6 to typically have values between 1 and 0.1 Hz [34]. The light travel time between LIGO detectors is 10 ms, leading to a coincident glitch rate of $R_{\rm gl} \sim 1$ Hz $\times$ 1 Hz $\times$ 0.01 s $= 0.01$ Hz.

False alarm rates are estimated by analyzing time-shifted data, or "time slides." If the data from one detector is shifted by more than the light travel time to another detector, there will be no coincident gravitational wave signals. Because the rate of glitches completely dominates the rate of GW signals, analyzing time-shifted data all but guarantees that any coincidences are due to noise artifacts.

Consider the last quarter of LIGO's sixth science run (S6D) which lasted for $\sim$ 50 days. A so-called "three sigma" detection requires an event more significant than any background coincidences found in ~ 300 time slides. The background estimate from 300 time slides corresponds to 40 years of data, and $N_{\rm gl} \sim 1 \times 10^7$. Equation 14 predicts that events with $\ln \mathcal{B}_{\mathcal{S},\mathcal{G}} \gtrsim 16$ would be detected with better than three-sigma confidence.

To test this prediction we compute the Bayes factors for the coincident events in time slides of the S6D data found by the `coherent WaveBurst` algorithm [17]. Figure 3 shows the cumulative glitch rate as a function of $\ln \mathcal{B}_{\mathcal{S},\mathcal{G}}$ i.e. the y-axis is the rate at which coincident glitches were found with Bayes factors greater than the corresponding value on the x-axis. The distribution steeply decreases with increasing Bayes factor, and does not show evidence of leveling-off with a broad "tail" in the background that has limited previous searches. See Ref. [18] for a detailed study of how `BayesWave` can improve detection confidence of existing burst searches. Furthermore, the distribution ends at $\ln \mathcal{B}_{\mathcal{S},\mathcal{G}} \sim 15$ which is consistent with our analytic prediction for the background. Ultimately we should be able to turn arguments about the expected background rate into a prior odds ratio between the glitch and signal model. For the immediate future we elect to take a more conservative approach and continue using background studies to estimate the false alarm rate and therefore the detection significance. There is no guarantee that the non-Gaussian noise in future GW data will bear any resemblance to what was found during S6.

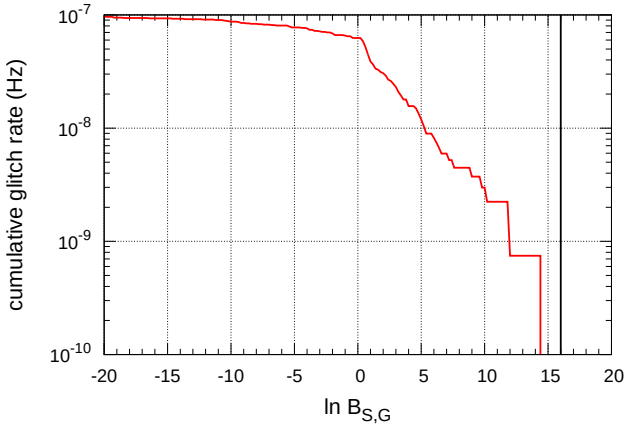

FIG. 3. The cumulative rate of glitches as a function of Bayes factor from time-slide studies using `BayesWave` and the S6D data set. The black, vertical line shows the expected value for the most significant event using the limit in Equation 14. Because this represents 300 time-slides of the data set, we see that the sine-gaussian injections above network SNR 10 are detected with marginal significance, whereas many WNB signals above network SNR 25 were "gold-plated" detections. Our findings are in excellent agreement with the analytic approximations in this work.

Comparing our inferred background rate to Figure 2 we find that sine-Gaussian waveforms in a two detector network will be detected at false alarm rates that suggest marginal significance at any reasonable SNR, similar to performance seen in past burst searches. However, unlike previous burst searches, we find that IMBH and white noise bursts are detectable with very high significance. Figure 5 in the Appendix shows the number of wavelets used to recover each waveform morphology as a function of injected SNR and provides supporting evidence that waveforms that require more wavelets typically provide higher Bayes factors.

What is required for a high confidence, or "five-sigma," detection? For this case, we seek a p-value of less than 3×10^{-7}, and so demand our event be louder than the loudest event in 3×10^6 time slides. For S6D this leads to $N_{\rm gl} \sim 10^{11}$, and an expected loudest event $\langle \ln \mathcal{B}_{\mathcal{S},\mathcal{G}} \rangle_{\rm background} \sim 25$. We have already seen that single wavelet events can not reach this level at any reasonable SNR but applying the scaling law in Equation 12, we find that such a "gold-plated" detection could be achieved at reasonable SNR with as few as two or three wavelets. For example, the IMBH and white noise burst signals in Figure 2 added to the same data we used to estimate the background by far exceed the Bayes factor which corresponds to a false alarm probability of $\sim 3 \times 10^{-7}$. This is an important feature of the `BayesWave` pipeline: *Gold-plated detections of short-duration signals are possible even in the presence of a significant glitch population.*

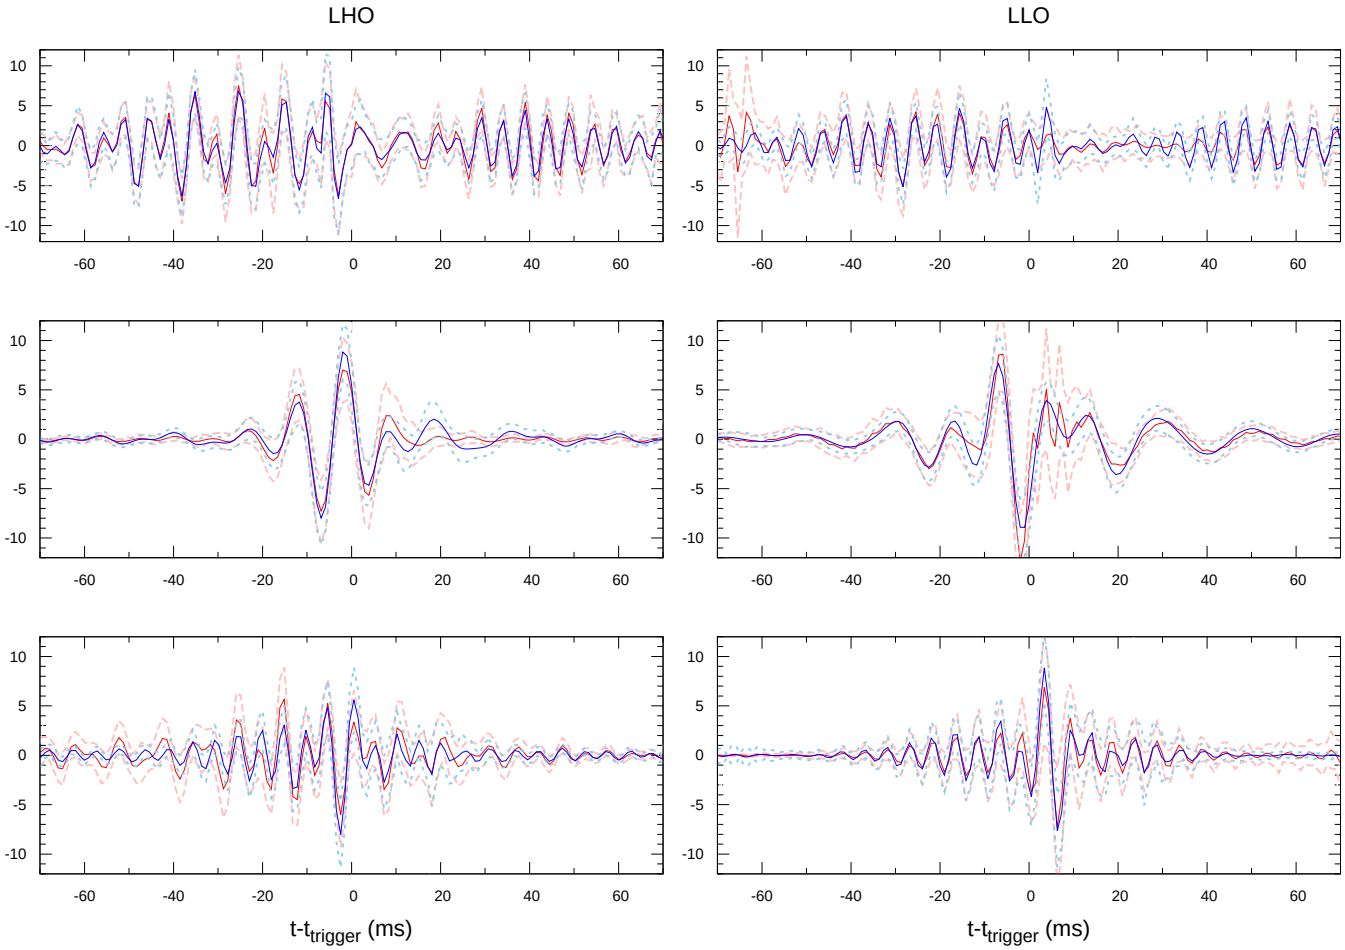

FIG. 4. Reconstructed whitened, time-domain signal- and glitch-model waveforms for S6D background events. Solid (red/blue) line is the median (glitch/signal) waveform. Dashed lines of corresponding color show the 2σ errors on the reconstructed waveforms. Each row shows one of the three most significant background events. Left column is the waveform in Hanford. Right column is Livingston. From top to bottom $\ln \mathcal{B}_{\mathcal{S},\mathcal{G}}$ was $[15, 12, 12]$. The overlap between the glitch model and the signal model was $[91\%, 93\%,$ and $86\%]$, respectively.

V. DISCUSSION

In this paper we have demonstrated `BayesWave`'s utility as a follow-up analysis for GW burst searches. By analyzing data from the sixth LIGO science run (S6) which took place from 2009-2010 we have shown that high confidence detections are achievable using `BayesWave` as a follow-up analysis despite the high rate of noise transients in the data. When used to follow-up short-duration gravitational wave triggers, `BayesWave` has been shown to significantly reduce the rate of false-alarms while remaining sensitive to a wide range of signals [18]. For insight into how `BayesWave` takes advantage of Bayesian model selection to separate signals and glitches we presented an analytical framework and found simple expressions which provide approximations to our full numerical analysis on real data. The results show that `BayesWave` has several novel features, when compared with other Burst pipelines:

- The detection statistic directly compares the evidence for an astrophysical signal with a glitch model, as opposed to calculating a likelihood derived from Gaussian noise.
- `BayesWave` places emphasis on the time-frequency complexity and network coherence of an event, rather than just its strength, to distinguish signals from glitches
- The background distribution shows no evidence of "tails" at high values.

In order to emphasize the importance of including the glitch model in a statistical framework, best fit waveforms for the signal and glitch models for the most "signal-like" background events in S6D are shown in Figure 4. For these examples of real glitches, the signal and glitch

model are shown to very nearly agree. Because glitches can be so successful in imitating real gravitational wave signals, pipelines which attempt to reject these events with tunings and cuts face a major challenge. Instead, `BayesWave` attempts to accurately assess the probability of such coincident glitches arising from chance. This approach places a lower weight to events with simple time-frequency structure that could plausibly arise simultaneously in two or more instruments, regardless of their SNR.

The detection statistic described in this work, $\mathcal{B}_{\mathcal{S},\mathcal{G}}$, represents the likelihood ratio for two competing models: the data contains a glitch, or the data contains an astrophysical signal. The purist may object to this application of the Bayes factor, instead favoring the Bayesian odds ratio between the signal and glitch model. The prior odds ratio between these models is the ratio of the expected coincident and coherent glitch rate to the expected rate of GW signals. While the rate of GW signals is unknown, we have shown that the measured background distribution is consistent with our analytic predictions using the LIGO glitch rate. This consistency suggests that the `BayesWave` model is a good fit to actual LIGO data and Bayes factors calculated by `BayesWave` will serve as a robust means for correctly identifying signals and glitches. In principle, glitches with non-flat distributions in f and Q, especially if similarly distributed in multiple detectors, could invalidate this agreement. Should that be the case, the posterior distribution of background events can easily be folded in to our analysis as a prior on the glitch model. Because the glitch population in earlier LIGO data will likely differ from that of the advanced detectors, we will continue to rely on the brute-force approach of using time slides to estimate the significance of a candidate event and use what is learned to further improve our priors for subsequently collected data.

As the capabilities of ground based detectors continues to improve so too must our analysis. The work presented here represents a snap shot of `BayesWave`'s capabilities as the algorithm continues to advance. Further development is underway to relax the requirement of elliptical polarization for the signal model (improving the detection efficiency for unpolarized signals) and to account for glitches and signals appearing in the same segment of data (reducing false dismissals due to near-coincidence with glitches). Nonetheless, based on the thorough performance studies in real LIGO data reported in this work we conclude that `BayesWave` is prepared to decisively aid in the detection and characterization of GW bursts in the advanced detector era.

VI. ACKNOWLEDGMENTS

We acknowledge numerous discussions with Reed Essick, Erik Katsavounidis, and Salvatore Vitale which helped motivate the detailed analytic derivation of the evidence in the Appendix; James Clark for thorough comments and suggestions on an earlier draft of this paper; Sergey Klimenko and Francesco Salemi for help finding and understanding the cWB output files; and Kent Blackburn, Vicky Kalogera, Tjonnie Li, Patricia Schmidt, and Michele Vallisneri for helpful conversations about this work. TBL acknowledges the support of NSF LIGO grant, award PHY-1307020. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058. This paper carries LIGO Document Number LIGO-P1500083.

Appendix: Bayes Factors

The Laplace approximation for the evidence is given by

$$Z = p(d|\boldsymbol{\theta}_{\rm MAP},\mathcal{H})p(\boldsymbol{\theta}_{\rm MAP}|\mathcal{H})\,(2\pi)^{D/2}\sqrt{\det C_{\boldsymbol{\theta}}} \quad \text{(A.1)}$$

where $\boldsymbol{\theta}_{\rm MAP}$ are the maximum a posteriori parameters, D is the model dimensions, and $C_{\boldsymbol{\theta}}$ is the full parameter covariance matrix, which we can estimate from the inverse of the Fisher information matrix $\Gamma_{\boldsymbol{\theta}}$. If the prior is uniform for all parameters, the prior density is equal to the inverse of the prior volume: $p(\boldsymbol{\theta}_{\rm MAP}) = 1/V_{\rm prior}$. We recognize the collection of terms $(2\pi)^{D/2}\sqrt{\det C_{\boldsymbol{\theta}}}$ as the posterior volume. Thus, for uniform priors, the evidence is given by the product of the MAP likelihood times the ratio of the posterior to prior volume, which is referred to as the Occam penalty. In the case of `BayesWave` the priors on most parameters are flat, with the important exception of the amplitude or the signal-to-noise ratio (which depends of the amplitude, quality factor, central frequency and noise spectral density).

Dropping terms down by factors of e^{-Q^2} relative to leading order, the Fisher matrix for a single wavelet using the parameters $\{t_0, f_0, Q, \ln A, \phi_0\}$ is given by

$$\Gamma_{\boldsymbol{\lambda}} = \mathrm{SNR}^2 \begin{pmatrix} \frac{4\pi^2 f_0^2(1+Q^2)}{Q^2} & 0 & 0 & 0 & -2\pi f_0 \\ 0 & \frac{3+Q^2}{4f_0^2} & -\frac{3}{4Qf_0} & -\frac{1}{2f_0} & 0 \\ 0 & -\frac{3}{4Qf_0} & \frac{3}{4Q^2} & \frac{1}{2Q} & 0 \\ 0 & -\frac{1}{2f_0} & \frac{1}{2Q} & 1 & 0 \\ -2\pi f_0 & 0 & 0 & 0 & 1 \end{pmatrix} \quad \text{(A.2)}$$

The determinant of $\Gamma_{\boldsymbol{\lambda}}$ is

$$\det\Gamma_{\boldsymbol{\lambda}} = \frac{1}{\det C_{\boldsymbol{\lambda}}} = \frac{\pi^2\,\mathrm{SNR}^{10}}{2\,Q^2}\,. \quad \text{(A.3)}$$

The expectation value of the MAP log likelihood is given by (see page 31 of Ref. [30] and references therein)

$$\mathrm{E}[\ln p(d|\boldsymbol{\theta}_{\rm MAP})] = \text{const.} + \frac{D}{2} \quad \text{(A.4)}$$

The constant is independent of the signal model. The $D/2$ term comes from more complicated models being able to better fit features in the Gaussian noise.

Each wavelet is described by 5 parameters, and has

$$\sqrt{\det C_{\boldsymbol{\lambda}_i}} = \frac{\sqrt{2}Q_i}{\pi\,\mathrm{SNR}_i^5} \tag{A.5}$$

where $\mathrm{SNR_i}$ is the signal-to-noise ratio for wavelet i. Assuming that the N wavelets used in the reconstruction have little overlap with each other, the total posterior volume for the wavelet model is

$$\sqrt{\det C_{\boldsymbol{\lambda}}} = \prod_{i=1}^{N} \frac{\sqrt{2}Q_i}{\pi\,\mathrm{SNR}_i^5} \tag{A.6}$$

`BayesWave` has a non-trivial amplitude prior which needs to be taken into account. One choice would be a uniform in volume prior on the source distribution, which is equivalent to a prior on the distance D that scales as $p(D)\ D^2$. Since amplitude and distance are inversely related, we have $D^2 dD \sim A^{-4} dA \sim A^{-3} d\ln A$. Here we have made the change of variables to $\ln A$ since this is the parameter used to compute the Fisher matrix. This prior is improper, and to make it proper a minimum amplitude cut-off A_* (maximum distance) has to be introduced. The properly normalized uniform-in-volume prior is

$$p(\ln A) = 3\left(\frac{A_*}{A}\right)^3 . \tag{A.7}$$

An alternative approach, used by `BayesWave` in this work, is to adopt different physically motivated priors on the signals and glitches that are given as functions of the SNR. For glitches the SNR is given by

$$\mathrm{SNR} \simeq \frac{A\sqrt{Q}}{\sqrt{2\sqrt{2\pi} f_0 S_n(f_0)}}\,, \tag{A.8}$$

while for signals the SNR is given by the same expression, but with the individual detector noise spectral density replaced by the network average

$$\bar{S}_n(f_0) = \left(\sum_i \frac{F_{+,i}^2 + \epsilon^2 F_{\times,i}^2}{S_{n,i}(f_0)}\right)^{-1} . \tag{A.9}$$

Thus the signal-model SNR depends on $A, Q, f_0, \boldsymbol{\Omega}$. The priors on the signal model and glitch model are given in terms of a prior on the SNR, $p(\mathrm{SNR})$. Making the change of variables from SNR to $\ln A$ yields

$$p(\ln A|\mathcal{G}) = \left(\frac{\mathrm{SNR}}{\mathrm{SNR}_*}\right)^2 e^{-\mathrm{SNR}/\mathrm{SNR}_*} \tag{A.10}$$

for the glitch model and

$$p(\ln A|\mathcal{S}) = \frac{3}{4}\left(\frac{\mathrm{SNR}}{\mathrm{SNR}_*}\right)^2 \frac{1}{(1+\mathrm{SNR}/4\mathrm{SNR}_*)^5} \tag{A.11}$$

for the signal model.

If we further assume little correlation between the wavelet model and the Gaussian noise model, then the expected value for the log Bayes Factor between the glitch plus noise model and the noise model in a single detector is

$$\ln \mathcal{B}_{\mathcal{G},\mathcal{N}} = \frac{\mathrm{SNR}^2}{2} + \frac{5N}{2}\left(1+\ln(2\pi)\right) + \sum_{n=1}^{N} \ln\left(\frac{\sqrt{2}Q_n}{\pi\,\mathrm{SNR}_n^5}\right) + \ln p(\boldsymbol{\lambda}_{\mathrm{MAP}}|\mathcal{G})\,. \tag{A.12}$$

Here SNR^2 is the signal-to-noise ratio of the signal or glitch in that detector. Later when considering a network of detectors the SNR^2 will refer to the network signal-to-noise ratio of the signal.

If the wavelet model in one detector uses N wavelets, and assuming little overlap between wavelets, then

$$\mathrm{SNR}^2 = \sum_{n=1}^{N} \mathrm{SNR}_n^2 = N\,\overline{\mathrm{SNR}}^2 \tag{A.13}$$

Based on simulations, we find that the average number of wavelets used by `BayesWave` increases linearly with the total SNR, and that the average per wavelet $\overline{\mathrm{SNR}}$ increases with the total SNR as a waveform-dependent power law. Writing $N = 1 + \beta\mathrm{SNR}$ and $\overline{\mathrm{SNR}} = \alpha\,\mathrm{SNR}^a$, we find that the values of β, α, and a depend on the waveform morphology, with α and β increasing, and a decreasing, as the time-frequency structure of the waveform becomes more complicated. In the case of a constrained model using a fixed number of wavelets the average SNR per wavelet always increases linearly with the total SNR, though with a proportionality less than one for anything other than sine-gaussians.

Fig. 5 shows the average number of wavelets (left panel) and SNR per wavelet (right panel) as a function of SNR for the three different waveform morphologies studied in this paper–sine Gaussians, binary black hole mergers, and white noise bursts–added to simulated Gaussian noise from a single detector at Advanced LIGO sensitivity. Each simulation was repeated for several Gaussian noise realizations. Plotted are the average and one standard deviations of the mean, plus lines that show the scaling relations using the best fit values for $\{\alpha, \beta, a\}$.

Starting with a simplified model of two aligned collocated detectors, the signal model does not need any

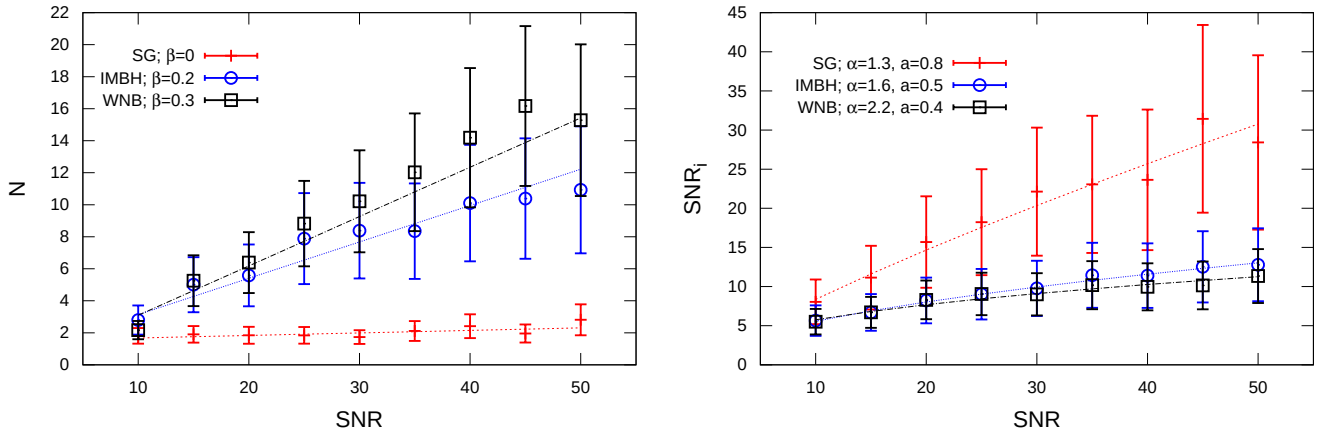

FIG. 5. Number of wavelets (left) and SNR per wavelet (right) found for simulated GWs as a function of the signal's SNR. Simulated waveforms include sine Gaussians (red), intermediate mass binary black holes (blue), and white noise bursts (black). The GW signals were added to simulated Gaussian noise at full design sensitivity for a single Advanced LIGO detector averaged over several Gaussian noise realizations. Error bars represent one standard deviation of the mean. Over-plotted are SNR scalings found in the text using the best fit values β, α, and a.

extrinsic parameters and the log Bayes factors are

$$\ln \mathcal{B}_{\mathcal{S},\mathcal{G}} = \left(\frac{5N^{\mathcal{S}}}{2} - \frac{5N^{\mathcal{G}}}{2}\right)(1+\ln(2\pi)) + \sum_{n=1}^{N^{\mathcal{S}}} \ln\left(\frac{\sqrt{2}Q_n}{\pi\,\mathrm{SNR}_n^5}\right) - \sum_{n=1}^{N^{\mathcal{G}}} \ln\left(\frac{\sqrt{2}Q_n}{\pi\,\mathrm{SNR}_n^5}\right) + \ln p(\boldsymbol{\lambda}_{\mathrm{MAP}}|\mathcal{S}) - \ln p(\boldsymbol{\lambda}_{\mathrm{MAP}}|\mathcal{G}). \quad \text{(A.14)}$$

If we assume that all the wavelets have the same quality factor Q and individual signal-to-noise ratios SNR_i, then the individual SNRs in each detector are $\mathrm{SNR}/\sqrt{2}$. Thus for the glitch model $\mathrm{SNR}_i^2 = \alpha\,\mathrm{SNR}/\sqrt{2}$, while for the signal model $\mathrm{SNR}_i = \alpha\mathrm{SNR}^a$. Then the Bayes Factor between the signal and glitch models for two collocated detectors is

$$\ln \mathcal{B}_{\mathcal{S},\mathcal{G}} = \frac{5N^{\mathcal{S}}}{2} - \frac{5N^{\mathcal{G}}}{2} + N^{\mathcal{S}} \ln\left(\frac{4\pi^{1/2}Q}{\mathrm{TF}\,\Delta Q(\alpha\mathrm{SNR}^a)^5}\right) - N^{\mathcal{G}} \ln\left(\frac{4\pi^{1/2}Q}{\mathrm{TF}\,\Delta Q(\alpha\mathrm{SNR}^a/\sqrt{2}))^5}\right) + N^{\mathcal{S}} \ln p(\ln A|\mathcal{S}) - N^{\mathcal{G}} \ln p(\ln A|\mathcal{G}) \quad \text{(A.15)}$$

Here TF is the time frequency volume, and ΔQ is the prior range for Q. Note that the contributions from the amplitude prior introduce important SNR scalings into the Bayes factor. For a single sine-Gaussian model the posterior volume terms introduce a ln(SNR) scaling to the log Bayes factor. For `BayesWave` the scaling with SNR is far more complicated. On complex waveforms both $N^{\mathcal{S}}$ and $N^{\mathcal{G}}$ scale linearly with the SNR, so the posterior volume introduces terms that scale as $\mathrm{SNR}\ln(\mathrm{SNR})$. The amplitude prior for the signal model introduces a similar $\mathrm{SNR}\ln(\mathrm{SNR})$ scaling, along with a more complicated scaling of the form $\mathrm{SNR}\ln(1+\mathrm{SNR}/4\mathrm{SNR}_*)$. The amplitude prior for the glitch model introduces $\mathrm{SNR}\ln(\mathrm{SNR})$ terms, in addition to a term that scales as SNR^2, though this term does not start to dominate until very high SNRs ($\mathrm{SNR} > 50$ for typical choices of SNR_*). In the fixed dimension case the `BayesWave` scaling is dominantly of the form ln(SNR) for moderate SNRs. At very high SNRs the quadratic dependence of the full `BayesWave` scaling is replaced by a linear scaling in SNR.

There are several assumptions that went into the derivation of the signal-to-glitch Bayes Factor for `BayesWave` that are rather crude. The worst approximations are that the wavelets used in the reconstruction all have roughly the same Q and signal-to-noise ratio.

While on average the scaling $\mathrm{SNR}_i = \alpha \mathrm{SNR}^a$ is quite robust, the SNR_i for individual wavelets never go much below the value set by the peak of the SNR prior, SNR_* so that $\mathrm{SNR}_i \gtrsim \mathrm{SNR}_*$. This means that the linear scaling typically only holds for network SNRs greater than around 10 or 12. Rather than assuming the same quality factor for each wavelet we could use the average value. For Q distributed uniformly in the range $Q \in [Q_1, Q_2]$ we have

$$\mathrm{E}\left[\ln\left(\frac{\Delta Q}{Q}\right)\right] = 1 + \frac{Q_2 \ln(\Delta Q/Q_2) - Q_1 \ln(\Delta Q/Q_1)}{\Delta Q}, \tag{A.16}$$

and

$$\mathrm{Var}\left[\ln\left(\frac{\Delta Q}{Q}\right)\right] = \frac{Q_1 Q_2 (\ln(\Delta Q/Q_2) - \ln(\Delta Q/Q_1))^2}{\Delta Q^2} - 1, \tag{A.17}$$

1. Extrinsic Parameters

Implicit in the preceding derivation was the assumption that the overlap of any two wavelets $(\Psi_i|\Psi_j) \sim 0$ and, as a consequence, the parameter correlation matrix for the wavelets was block-diagonal. This assumption is reasonable since each wavelet collects the power in a certain time-frequency volume disfavoring highly overlapping wavelets. For mis-aligned detectors we can write the parameter correlation matrix for the signal model in block-form as

$$C_{\boldsymbol{\theta}} = \begin{pmatrix} C_{\boldsymbol{\lambda}} & C_X \\ C_X^T & C_{\boldsymbol{\Omega}} \end{pmatrix} \tag{A.18}$$

where $C_{\boldsymbol{\lambda}}$ is the $5N \times 5N$ correlation matrix for the intrinsic wavelet parameters, $C_{\boldsymbol{\Omega}}$ is the 4×4 correlation matrix for the extrinsic parameters and C_X is the $5N \times 4$ cross-correlation matrix that mixes the extrinsic and intrinsic parameters. The Fisher matrix $\Gamma_{\boldsymbol{\theta}} = C_{\boldsymbol{\theta}}^{-1}$ can similarly be decomposed:

$$\Gamma_{\boldsymbol{\theta}} = \begin{pmatrix} \Gamma_{\boldsymbol{\lambda}} & \Gamma_X \\ \Gamma_X^T & \Gamma_{\boldsymbol{\Omega}} \end{pmatrix} \tag{A.19}$$

Now, for partitioned symmetric matrices we have (see page 46 of the Matrix Cookbook [35])

$$\det C_{\boldsymbol{\theta}} = \frac{\det C_{\boldsymbol{\Omega}}}{\det \Gamma_{\boldsymbol{\lambda}}}, \tag{A.20}$$

which implies that the volume of the posterior factors into extrinsic and intrinsic pieces, where the intrinsic part has exactly the same form as for the glitch model:

$$\begin{aligned} V_{\mathcal{S}} &= (2\pi)^{D/2} \sqrt{\det C_{\boldsymbol{\theta}}} \\ &= (2\pi)^{5N^{\mathcal{S}}/2+2} \sqrt{\det C_{\boldsymbol{\Omega}}} \prod_{i=1}^{N^{\mathcal{S}}} \left(\frac{\sqrt{2} Q_i}{\pi \, \mathrm{SNR}_i^5}\right) \end{aligned} \tag{A.21}$$

Putting all the pieces together we have

$$\begin{aligned} \ln \mathcal{B}_{\mathcal{S},\mathcal{G}} = &\left(\frac{5N^{\mathcal{S}}}{2} + 2 - \frac{5N^{\mathcal{G}}}{2}\right)(1 + \ln(2\pi)) + \sum_{n=1}^{N^{\mathcal{S}}} \ln\left(\frac{\sqrt{2} Q_n}{\pi \, \mathrm{SNR}_n^5}\right) - \ln p(\boldsymbol{\lambda}_{\mathrm{MAP}}|\mathcal{G}) \\ &+ \ln\left(\frac{\sqrt{\det C_{\boldsymbol{\Omega}}}}{4\pi^2}\right) - \sum_{n=1}^{N^{\mathcal{G}}} \ln\left(\frac{\sqrt{2} Q_n}{\pi \, \mathrm{SNR}_n^5}\right) + \ln p(\boldsymbol{\lambda}_{\mathrm{MAP}}|\mathcal{S}) \end{aligned} \tag{A.22}$$

From here we can insert the SNR scalings for $N^{\mathcal{S}}$, $N^{\mathcal{G}}$ and the SNR_n and include the explicit expression for the intrinsic parameter volumes in an effort to make quantitative predictions. While the expressions are more complicated than the aligned collocated case the scalings with SNR are the same.

We are unable to derive an analytic expression for $\sqrt{\det C_{\boldsymbol{\Omega}}}$. Additionally, there is the problem that some of the extrinsic parameters, most notably the ellipticity and polarization angle, are poorly constrained and Fisher matrix estimates are unreliable. As a result the posterior volume does not scale as SNR^{-4} as we naively expect from the Fisher matrix, but as some lower power such as SNR^{-2} or SNR^{-3}. One way to incorporate the restriction that the posterior not exceed the prior is to elevate the extrinsic parameters from their fundamental domain (with periodic boundary conditions) to the universal cover, and introduce a Gaussian prior on the parameters that restricts the posterior volume to be no larger than the prior volume. The negative Hessian of second derivatives of the log of this prior is added to the Fisher matrix (so that the Fisher matrix describes the curvature of the posterior, not just the likelihood). Numerically computing the posterior volume as a function of SNR using this approach shows that the posterior volume scales as $\mathrm{SNR}^{-\gamma}$, where the exponent γ is weakly dependent on the SNR, varying between roughly 2 and 3 across the range of SNRs we expects to encounter, as shown in Fig. 5.

The end result is that including the intrinsic parameters increases the dimension of the signal model by between 2 and 3 degrees of freedom, not 4 as we would naively expect. Thus, the overall scaling for the single sine-Gaussian Bayes factor should scales as $\ln \mathcal{B}_{\mathcal{S},\mathcal{G}} \sim$

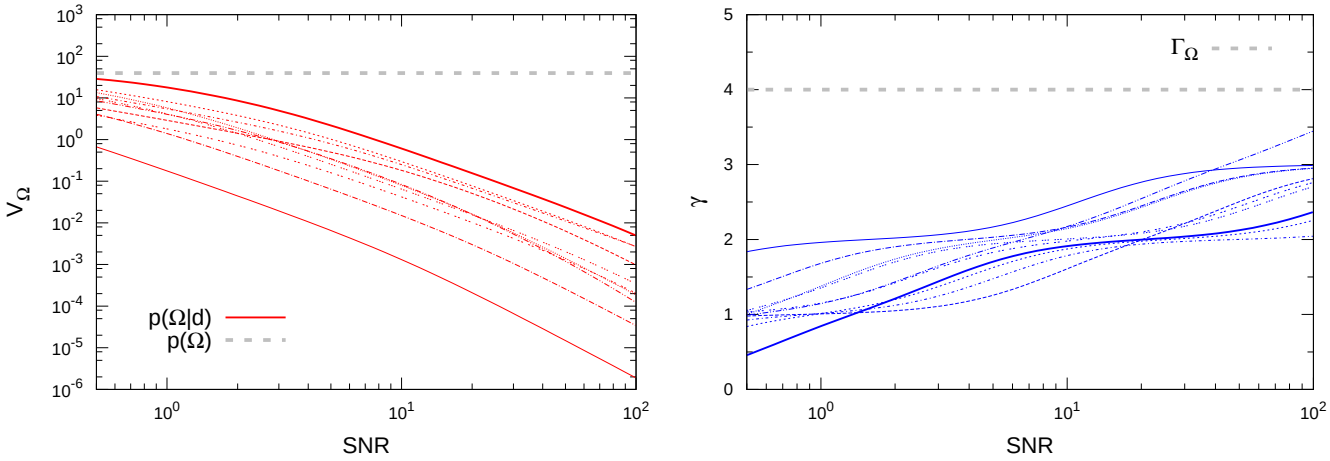

FIG. 6. The panel on the left shows the scaling of the extrinsic posterior volume with SNR for ten randomly chosen single sine-Gaussians in a two detector network. The panel on the right shows the SNR dependence of the slope parameter γ.

$(5 \rightarrow 6)\ln \mathrm{SNR}$. The scaling for more elaborate waveforms is far more complicated. Ultimately it is this added complexity that enables `BayesWave` to assign high-confidence to detection candidates of non-trivial GW signals.

[1] B. Abbott *et al.* (LIGO Scientific Collaboration), Rept. Prog. Phys. **72**, 076901 (2009), arXiv:0711.3041 [gr-qc].
[2] T. Accadia *et al.*, Journal of Instrumentation **7**, P03012 (2012).
[3] B. Abbott *et al.* (LIGO Scientific Collaboration), Rept.Prog.Phys. **72**, 076901 (2009), arXiv:0711.3041 [gr-qc].
[4] L. Blackburn, L. Cadonati, S. Caride, S. Caudill, S. Chatterji, *et al.*, Class.Quant.Grav. **25**, 184004 (2008), arXiv:0804.0800 [gr-qc].
[5] J. Abadie *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. D **82**, 102001 (2010), arXiv:1005.4655 [gr-qc].
[6] J. Aasi *et al.* (VIRGO Collaboration), Class.Quant.Grav. **29**, 155002 (2012), arXiv:1203.5613 [gr-qc].
[7] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **80**, 102002 (2009), arXiv:0904.4910 [gr-qc].
[8] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys.Rev. **D80**, 102001 (2009), arXiv:0905.0020 [gr-qc].
[9] J. Abadie *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. D **85**, 122007 (2012), arXiv:1202.2788 [gr-qc].
[10] E. Thrane and M. Coughlin, Phys. Rev. Lett. **115**, 181102 (2015), arXiv:1507.00537 [astro-ph.IM].
[11] J. Veitch, V. Raymond, B. Farr, W. Farr, P. Graff, *et al.*, Phys.Rev. **D91**, 042003 (2015), arXiv:1409.7215 [gr-qc].
[12] R. Essick, S. Vitale, E. Katsavounidis, G. Vedovato, and S. Klimenko, (2014), arXiv:1409.2435 [astro-ph.HE].
[13] R. Lynch, S. Vitale, R. Essick, E. Katsavounidis, and F. Robinet, "2015," In preperation.
[14] J. Veitch and A. Vecchio, Phys. Rev. D **81**, 062003 (2010), arXiv:0911.3820 [astro-ph.CO].
[15] N. J. Cornish and T. B. Littenberg, Class. Quant. Grav. **32**, 135012 (2015), arXiv:1410.3835 [gr-qc].
[16] T. B. Littenberg and N. J. Cornish, Phys.Rev. **D91**, 084034 (2015), arXiv:1410.3852 [gr-qc].
[17] S. Klimenko, I. Yakushin, A. Mercer, and G. Mitselmakher, Class. Quantum Grav. **25**, 114029 (2008), arXiv:arXiv:0802.3232 [gr-qc].
[18] J. B. Kanner, T. B. Littenberg, N. Cornish, M. Millhouse, E. Xhakaj, F. Salemi, M. Drago, G. Vedovato, and S. Klimenko, ArXiv e-prints (2015), arXiv:1509.06423 [astro-ph.IM].
[19] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis, *8th Gravitational Wave Data Analysis Workshop (GWDAW 2003) Milwaukee, Wisconsin, December 17-20, 2003*, Class. Quant. Grav. **21**, S1809 (2004), arXiv:gr-qc/0412119 [gr-qc].
[20] S. Chatterji, *The search for gravitational wave bursts in data from the second LIGO science run*, Ph.D. thesis, Massachusettes Institute of Technology, http://hdl.handle.net/1721.1/34388 (2005).
[21] S. Chatterji, A. Lazzarini, L. Stein, P. J. Sutton, A. Searle, *et al.*, Phys.Rev. **D74**, 082005 (2006), arXiv:gr-qc/0605002 [gr-qc].
[22] P. J. Sutton *et al.*, New J. Phys. **12**, 053034 (2010), arXiv:0908.3665 [gr-qc].
[23] J. Abadie *et al.* (LIGO), Phys. Rev. **D83**, 042001 (2011), arXiv:1011.1357 [gr-qc].
[24] D. Gamerman, *Markov Chain Monte Carlo: Stocastic Simulation of Bayesian Inference* (Chapman and Hall, London, 1997).
[25] R. Swendsen and J. Wang, Phys. Rev. Lett. **57**, 2607 (1986).
[26] P. Goggans and Y. Chi, *Bayesian Inference and Methods in Science and Engineering*, edited by J. Rychert, G. Er-

ickson, and C. Smith (American Institute of Physics, USA, 2004).
[27] W. Vousden, W. M. Farr, and I. Mandel, ArXiv e-prints (2015), arXiv:1501.05823 [astro-ph.IM].
[28] P. Green, *Highly Structured Stochastic Systems*, edited by P. Green, N. Hjort, and S. Richardson (Oxford University Press, 2003).
[29] W. Anderson, P. Brady, D. Chin, J. Creighton, K. Riles, and J. Whelan, "Beam pattern response functions and times of arrival for earthbound interferometer," LIGO-T010110-00-Z (2001).
[30] C. Röver, *Bayesian inference on astrophysical binary inspirals based on gravitational-wave measurements.*, Ph.D. thesis, The University of Auckland, URL:http://hdl.handle.net/2292/2356 (2007).
[31] N. J. Cornish and T. B. Littenberg, Phys. Rev. D **76**, 083006 (2007), arXiv:arXiv:0704.1808 [gr-qc].
[32] M. Vallisneri, Phys. Rev. D **77**, 042001 (2008), arXiv:gr-qc/0703086.
[33] A. Buonanno *et al.*, Phys. Rev. D **76**, 104049 (2007), arXiv:0706.3732 [gr-qc].
[34] J. Aasi *et al.* (LIGO Scientific, VIRGO), Class.Quant.Grav. **32**, 115012 (2015), arXiv:1410.7764 [gr-qc].
[35] K. B. Petersen and M. S. Pedersen, *The Matrix Cookbook* (Technical University of Denmark, URL:http://www2.imm.dtu.dk/pubdb/p.php?3274, 2012).