Robust detection and characterization of generic gravitational wave transients

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

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

With the advanced LIGO and Virgo detectors nearing completion the detection of gravitational waves is imminent. To fully realize the potential for gravitational wave astronomy we must characterize the signals and infer the nature of the emitting objects. Extracting astrophysical information from gravitational wave detections is straightforward and well studied when detailed models for the waveforms are available, However, an important source of motivation for the field of gravitational wave astronomy is the potential for new discoveries. Recognizing unanticipated signals requires sophisticated data analysis techniques which do not depend on theoretical models. How to draw robust inferences from un-modeled sources of gravitational waves is not yet established, and detection methods are 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 showcase `BayesWave`'s capabilities by demonstrating glitch rejection, signal identification, and waveform characterization using data from LIGO and Virgo collected from 2009 to 2010.

I. INTRODUCTION

When the LIGO [1] and Virgo [2] observatories make their first detection of gravitational waves (GW) it will be a monumental achievement in the history of science. Detection itself is only part of the challenge. Gravitational wave astronomy comes to fruition through characterizing signals and using them to draw astrophysical inferences. Procedures for mining through the LIGO/Virgo data *searching* for signals and estimating the background have largely been in place since the initial interferometers were in operation. As detector sensitivity has improved through hardware upgrades [3, 4], and with it the prospects for detections, development of methods for *characterizing* candidate detections have become a priority. For template-based analyses hunting the mergers of compact objects (neutron stars or black holes) Bayesian stochastic sampling methods have been studied for over a decade, starting with the pioneering efforts of Chrstensen and Meyer [5], and culminating in `LALInference` – a robust data analysis pipeline for parameter estimation of compact binaries [6]. Search methods for transient un-modeled gravitational wave signals, or "bursts," have similarly been in place since the first observing runs with initial LIGO thanks to the stalwart burst detection pipeline, `coherent WaveBurst` [7]. Unlike the template-based analyses, characterizing un-modeled signals is a less well-defined problem and a flagship burst parameter estimation and model selection analysis has not been identified.

Attempts at making astrophysical inferences from the data without good models for the gravitational wave signal have focused on waveform and position reconstruction. Time-of-arrival differences between detectors [8, 9], and the relative amplitudes of the GW at each interferometer [10] provides sky-location information, with timing and amplitude uncertainty propagating through the calculations to estimate the error-bars on position reconstruction. The `LALInference` pipeline has been applied to the burst problem using a sine-Gaussian template as a surrogate waveform, the byproduct of which are robust Bayesian posteriors for the sky-location [11]. These position reconstruction methods provide useful information for electromagnetic follow-up observations of gravitational wave detections but do not yield sufficient (or any) insight into the gravitational wave signal itself.

Burst analyses targeting specific signal types can bear some insight into the gravitational wave source. Core-collapse supernova are important sources for Advanced LIGO and fall into the "burst" analyses because our theoretical understanding of the explosion mechanism, and resultant gravitational radiation, is not sufficient to use matched filtering methods. Numerical simulations do provide some information about possible waveform morphologies which can be harnessed for parameter estimation purposes. Summerscales et al used a Maximum Entropy method for detection, waveform reconstruction, and rudimentary parameter estimation by correlating the inferred signal with a catalog of waveforms from numerical simulations of core collapse supernovae [12]. These studies showed potential for distinguishing between models for density, angular momentum, differential rotation, and progenitor mass.

Methods using principal component analyses (PCA) where the principal components are constructed from catalogs of core-collapse supernova simulation waveforms

have recently come into vogue. Rover et al have demonstrated how progenitor mass, angular momentum, and equation of state parameters can be extracted from rotating stellar core collapse models [13]. Logue et al and Edwards et al differentiate between progenitor models by using PCAs constructed from simulations of the different catalogs as competing models for the data, then comparing the Bayesian evidence for each PC family [14, 15].

Still missing is a burst analysis which provides robust parameter estimation and waveform reconstruction for generic transient gravitational wave signals. Furthermore, previous searches for gravitational waves have been limited by transient, non-Gaussian noise artifacts, or glitches, which cause the various detection statistics to have heavier tails than they would for Gaussian noise [16–18]. Glitches are especially problematic for burst searches which, unlike binary in-spiral analyses, do not have the advantage of strong constraints on the GW signal morphology. A data analysis procedure which can provide robust parameter estimation, signal reconstruction, and glitch rejection for candidate generic gravitational wave transients is necessary to maximize the potential of gravitational wave observatories.

Recently we proposed `BayesWave` – a Bayesian algorithm to fill the gap in burst analysis by following-up candidate gravitational wave transient events found by the search pipelines, segregate signals from glitches, and provide robust signal characterization for arbitrary burst waveforms [19]. `BayesWave` compares the Bayesian evidence for the data containing a GW signal, an instrument artifact, or merely statistical fluctuations of the instrument background consistent with Gaussian noise. In the event that the candidate is of astrophysical origin, `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` also provides a complete characterization of the instrument noise including spectral estimation for the background Gaussian noise and "glitches" 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` [20]. 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 section II we will briefly re-introduce the `BayesWave` and `BayesLine` algorithms. Section III contains results showing how `BayesWave` successfully rejects glitches in LIGO data (III A) while still efficiently detecting simulated signals (III B) by using Bayesian evidence and a flexible, parameterized, model for the instrument noise. We also demonstrate how `BayesWave` can be used to inform our physical interpretation of detected signals without relying on theoretical models of burst waveforms in Section III C. Section IV provides a summary of our findings and discussion of future improvements for `BayesWave`.

II. METHOD

Fundamentally, we have constructed a complete parameterized model for the LIGO/Virgo detector output including noise and GW signals. `BayesWave` then uses a Markov chain Monte Carlo (MCMC) algorithm to produce samples for the posterior distribution function of the model parameters [21]. The model has three distinct components: A gravitational wave signal h that is elliptically polarized and is coherent across the network of detectors; glitches g – non-Gaussian noise transients that are independent in each interferometer; and stationary Gaussian noise which is fully characterized by its power spectral density $S_n(f)$, proportional to the variance of the noise at each frequency.

A. Gaussian noise model

Instrument noise n from gravitational wave detectors is well approximated by a zero mean, stationary, Gaussian distribution which is completely characterized by the frequency-dependent variance $\langle|\tilde{n}(f)|^2\rangle = \frac{T}{2}S_n(f)$. Here a tilde denotes a Fourier transform, T is the observation time and $S_n(f)$ is the noise power spectral density (PSD). The LIGO/Virgo noise budget is broadly dominated, from low to high frequency, by seismic noise, thermal noise from the mirror suspensions and coatings dominate, and quantum (photon shot) noise. Additional narrow band spectral lines which originate from a variety of sources including the mirror suspensions, the AC electrical supply, or sinusoidal motion imparted on the mirrors for calibration punctuate the otherwise smooth noise spectrum.

We have developed a two-component model to fit to the different qualitative features in the PSD. The broadband noise is parameterized by control points in frequency-PSD space which are interpolated between using a cubic spline fit. The spectral lines are modeled by a linear combination of Lorentzians. The number of spline control points and Lorentzians used in the model are determined using a trans-dimensional Reverse Jump Markov Chain Monte Carlo (RJMCMC) algorithm [27] . The Gaussian noise model uses a separate sampling algorithm, `BayesLine` to produce a posterior distribution for the instrument noise. Further details about the characteristics of the LIGO Gaussian noise, and the `BayesLine` algorithm, can be found in Ref [20].

B. Glitch and signal model

While stationary Gaussian noise is generally a good description for LIGO/Virgo noise the approximation breaks down with much higher regularity than the arrival of detectable gravitational waves. Any data analysis method must have a strategy for accounting for the non-stationary non-Gaussian noise. Most existing procedures opt for estimating the background rate of transient noise glitches and using that to adjust the threshold for detecting a gravitational wave signal. We instead approach the problem by developing a parameterized model glitches and incorporating it into our overall representation of the detector noise.

Because we don't know *a priori* the functional form of glitch or GW burst waveforms, our model for both must be flexible. Inspired by the work of Principe and Pinto [22, 23] which found that a modest number of wavelet basis functions were needed to accurately represent typical glitches in initial LIGO data, and our own experience of determining the number of white dwarf binaries which can be distinguished from a confusion-limited background in space-based gravitational wave detectors [24–26], we use a linear combination of Morlet-Gabor wavelets as our waveform model. Each basis function (wavelet) is described by five *intrinsic* parameters: central frequency f_0; central time t_0; amplitude A; quality factor Q; and phase offset ϕ_0; and is expressed in the time-domain as

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

where $\tau = Q/(2\pi f_0)$ and $\Delta t = t - t_0$. Like in `BayesLine`, the novel feature of `BayesWave` is that the *number* of wavelets included in the linear combination, N, is also a model parameter thereby requiring a RJMCMC to determine the number of wavelets needed for the model to be consistent with the data.

If the gravitational wave signal and the glitch model use the same functional form for the waveform, how can we distinguish the two? Glitches are not correlated between detectors so the linear combination of wavelets used to fit data from one observatory will be independent from the linear combination in the other. The signal is coherent across the network so we reconstruct the waveform as it would appear at the center of the Earth (geocenter) and forward-model the response of each detector with the addition of four *extrinsic* parameters: polar and azimuthal angles for the sky location (θ, ϕ); ellipticity ϵ to decompose the signal into the two gravitational wave polarizations + and ×; and the polarization angle ψ. Bayesian model selection is then used to determine which hypothesis – glitch or gravitational wave signal – is most consistent with the data. If a signal requiring $N_{\rm wavelet}$ wavelets to model deposits similar energy in each of the $N_{\rm detector}$ detectors the signal model will use $D_s = 5N_{\rm wavelet} + 4$ parameters while the glitch model, fitting each detector's data independently, will use $D_g = N_{\rm detector} \times 5N_{\rm wavelet}$ parameters. Occam's razor will favor the signal model. If a glitch appears in a single detector the signal model may be able to find extrinsic parameters which put zero signal energy in one detector, thereby achieving a similar fit to the data for the same number of wavelets, but will pay the price of carrying the extrinsic parameters, and Occam's razor will favor the glitch model.

Counting degrees of freedom is, of course, an overly simplistic way of hypothesis testing. Not all dimensions are created equally as unconstrained parameters should not penalize the model. The actual model comparisons are done by comparing the Bayesian evidence.

C. Computing the evidence

The central engine of `BayesWave` is a RJMCMC algorithm which, by construction, prefers the most parsimonious model for the data. RJMCMCs are used for model selection problems because they directly sample from the likelihood distribution in model-space. RJMCMCs require carefully tuned proposal distributions to adequately mix between models and can be notoriously difficult to implement, especially when the evidence ratio for the models under comparison has large dynamic range from one run to the next We have developed an RJMCMC implementation that mixes well within each model, sufficiently marginalizing over the number of basis functions needed for the glitch, signal, or noise model. Exchanges between models, on the other hand, are challenging and we currently can not rely on RJMCMC to select between, for example, the signal or glitch model. For model selection to determine if the data being analyzed contain a gravitational wave or not, we use a more brute-force method for evidence calculation – thermodynamic integration (TI) [28].

In our implementation of TI we use parallel tempering, where many Markov chains are run simultaneously, each sampling with a likelihood function $\mathcal{L}$ [25, 29]. The likelihood is modified by "temperature" T_i via $\mathcal{L}_i = \mathcal{L}^{1/T_i}$ where the subscript i is used to label the different Markov chains. Chains with $T \gg 1$ sample from the prior, while the chain with $T = 1$ samples the posterior. Chains can exchange parameters without violating detailed balance making parallel tempering a common approach to improve the efficiency and reliability of the MCMC. Thermodynamic integration uses the likelihood distribution of each chain to compute the log evidence by integrating the average log likelihood of each chain over the inverse temperature.

To summarize, we analyze data with a candidate event using three separate over-arching models (Gaussian noise, glitch, or gravitational wave signal). Within each analysis we use a parallel tempered RJMCMC to achieve the appropriate complexity, and produces model-averaged posterior distributions functions for parameters. Thermodynamic integration is then used to

compute the evidence for each model, which has been marginalized over the number of basis functions needed for each model. The entire procedure is discussed in detail in Refs. [19, 20].

III. RESULTS

In this section we will demonstrate BayesWave's glitch rejection, signal detection, and characterization capabilities. Model selection is performed by computing the Bayesian evidence for each model under consideration. Evidence ratios, or "Bayes factors," yield the relative likelihood for the two models. If we assume that prior preferences for different models are even the Bayes factor is equivalent to the odds ratio. The odds ratio tells us the relative probability that the two models are the correct description of the data. To simplify the interpretation of our results we will distill the Bayes factors into a decision criteria. BayesWave estimates the variance σ^2 of the statistical error in the evidence calculation and we will consider models that are preferred by greater than 3σ as "clear" preference for that model. [We need one more sentence to cover our embarrassment for using the Bayes Factor as a statistic. I can redo the background plots using 1,2,3 sigma contours. I could do the same for the efficiency curves but that would be a mess.]

A. Glitch rejection

The BayesWave algorithm's glitch-rejection performance is quantified by analyzing data from LIGO's sixth (S6) and Virgo's second (VSR2) science runs. The algorithm is tested on data with known problematic glitches by analyzing the most significant background events for the coherent WaveBurst burst search. To understand this study we must describe the procedure for estimating background. Gravitational wave detectors can not be shielded from incoming gravitational radiation so backgrounds must be estimated from data potentially containing gravitational wave signals. Because coincidence is required for detection (a signal in one interferometer is indistinguishable from a glitch) the background is estimated by time-shifting the data from one interferometer with respect to another to guarantee there are no coincident signals, and then running the time-shifted data through the detection pipeline. Events from the analysis of time-shifted data that score high in the ranking statistic – due to glitches in different detectors that fortuitously line up with one another – create the background distribution, and ultimately the thresholds for detection [30].

To demonstrate BayesWave's effectiveness we select the 100 highest-ranking background events with central frequency below 512 Hz from each 2-detector combination of S6/VSR2, making 300 background events total. BayesWaveis run on the time-shifted data as if it were coincident and asked to compare the evidence that the data contain a coherent gravitational wave signal, or incoherent but coincident glitches. Figure 1 shows the results of this test. The dependent variable ρ, which is related to the signal to noise ratio (SNR), is the ranking statistic used by cWB. The gray dashed histogram shows the number of background events at different values of ρ found by the original burst search which were then, in part, used to determine the threshold for detection. The red solid histogram shows the distribution of these preselected background events that are favored by the signal model over the glitch model by BayesWave. The blue dotted histogram applies an additional cut that the signal model is also preferred over the Gaussian noise model. By employing BayesWave to discriminate between signal and noise we see reduction in the background by almost an order of magnitude across ρ, completely eliminating all background events with $\rho \gtrsim 10$.

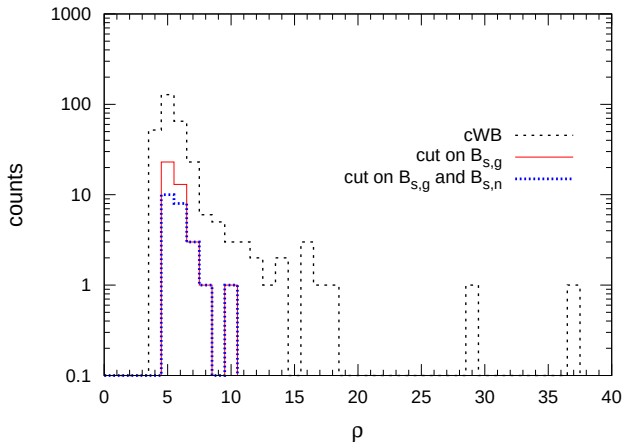

FIG. 1: Distribution of background events from the fourth quarter of S6/VSR2 run found by analyzing time shifted data (color online). The x-axis is ρ, the ranking statistic for the cWB search pipeline. The black, dashed histogram is the cWB background. The solid (red) histogram are the remaining background events after rejecting times when BayesWave preferred the glitch model over the signal. The dotted (blue) distribution uses an additional requirement that the signal model has to be favored over the Gaussian noise model. BayesWave provides dramatic reduction in the number of background events.

B. Signal detection

Any pipeline could be tuned to achieve similar glitch rejection capabilities to what is shown in Fig. 1 but there is a trade-off in detection efficiency. Rejecting background events is only useful if we do not mis-classify real signals as glitches. We demonstrate BayesWave's detection capabilities by adding simulated gravitational wave signals, "injections," to data taken during S6.

We inject two waveform morphologies: elliptically polarized sine-Gaussian waveforms (SG) and unpolarized

white-noise bursts (WNB). White noise bursts used in this analysis are simulated by band-passing white noise between 100 and 200 Hz and then time windowing the white noise with a Gaussian profile of duration 100ms. The h_+ and $h_\times$ polarizations are simulated independently resulting in unpolarized injections [31].

Because our basis functions for signal/glitch reconstruction are identical to the SG injections, analysis of these signals is equivalent to a template-based approach and provides important sanity checks for the pipeline. However, the difference in dimension between the glitch and the signal model for SG injection, which needs just a single basis function to reconstruct the waveform, is at a minimum. Form the point of view of segregating glitches and signals, we counterintuitively have the least fidelity for waveforms that match our basis functions.

WNB injections are, from the vantage point of waveform reconstruction, the worst-case-scenario for `BayesWave`. The injections will require a large number of basis functions to recover most of the power, and the injected waveforms are un-polarized whereas `BayesWave` assumes elliptical polarization.

Figure 2 summarizes the results from our study if signal injections. Each curve shows the efficiency with which `BayesWave` correctly identifies the data as containing a gravitational wave signal as a function of the signal to noise ratio. The left-hand panel has the efficiency curves for the signal model versus the Gaussian noise model plotted as a function of network SNR. We find the efficiency going to one for network SNRs around 10 for the sine-Gaussian waveforms (red, solid), and 20 for the white noise burst waveforms (blue, dashed). This is consistent with our expectations given our choice of signal model – for the SG injections `BayesWave` is a matched filtering search so we expect high efficiency. WNB signals, on the other hand, have more diffuse time frequency support, requiring more wavelets, and are not elliptically polarized, preventing `BayesWave` from ever achieving a perfect match to the signal.

The key piece of evidence that `BayesWave` does not sacrifice detection efficiency for the sake of glitch rejection comes from comparing the signal vs. glitch efficiency curves for the lowest SNR in the network (SG: magenta, dotted, WNB: cyan, dashed-dotted). Depending on the sky location and polarization of an injection, the signal may only be "detectable" in a single interferometer. A signal that only appears in one detector is indistinguishable from a glitch when we do not have strong constraints on the waveform morphology.

For sine-Gaussian waveforms, we find them to be completely distinguishable from Gaussian noise above SNRs of 10. For signal/glitch separation there are high network SNR injections that are identified as glitches – the network efficiency curve does not go to 1. At the point where lowest SNR in the network approaches 10 – i.e. the signal becomes detectable in each interferometer – we find 100% efficiency in discriminating these injections from glitches. High network SNR injections that are classified as glitches have unfortunate combinations of sky-location and polarizations such that all of the detectable power ends up in one interferometer. Without strong signal priors such events should be classified as glitches. None of the sine-Gaussian injections which we would expect to be detected – having measurable power in each interferometer – are misclassified as glitches.

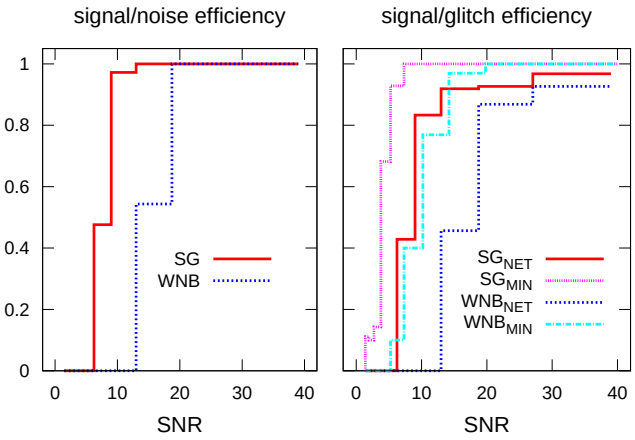

FIG. 2: Efficiency for selecting the signal model as a function of injected SNR compared to the Gaussian noise model (left) and glitch model (right). The signal-to-glitch panel has efficiency curves for the network SNR (SG: red/solid, WNB: blue/dashed) and for the lowest SNR in the network (SG: magenta/dotted, WNB: cyan/dash-dotted). All injections were performed in data from LIGO's sixth science run.

Similar conclusions can be drawn for the WNB injections despite the expected difficulty `BayesWave`'s signal model would have representing these waveforms with high fidelity. Despite the injected signal morphology being so different from the signal model basis functions, the only substantial difference between `BayesWave`'s performance on sine-Gaussian injections and white noise bursts is the injected signal to noise ratio where the waveforms become distinguishable from Gaussian noise. Similar to the SG injections, once there is detectable power in both interferometers injections are properly identified as signals.

C. Parameter estimation

[I'm out of gas. Anybody want to suggest a paragraph or two. Check out Figure 3].

IV. DISCUSSION

We have previously put forth a new way of characterizing un-modeled gravitational wave transient signals while simultaneously providing robust model selection comparisons between gravitational wave and glitchy noise hypotheses. In this paper we report on the first large-scale

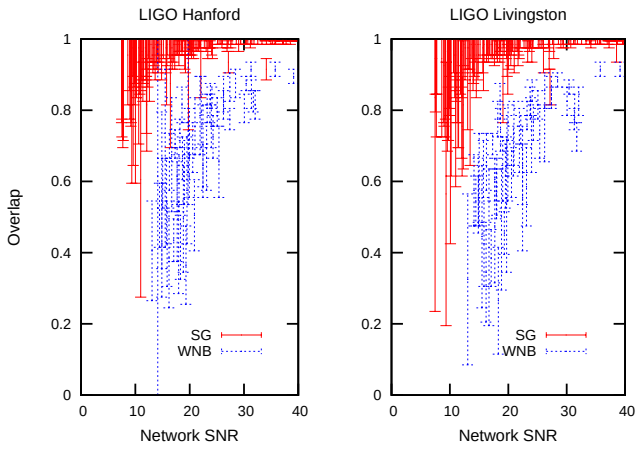

FIG. 3: Overlap of recovered waveform with injected waveform in each interferometer as a function of SNR. The solid (red) error bars represent the 90% credible interval of the overlap posterior distribution for detections of polarized sine-Gaussian waveforms. The blue (dotted) error bars are for the unpolarized white noise burst injections.

tests of BayesWave on data collected by the LIGO and Virgo observatories by analyzing times known to contain glitches which contributed to the search backgrounds, and by adding simulated GW signals to otherwise noise-only data. BayesWave has performed admirably at rejecting background events as being instrument noise without loss of detection efficiency.

BayesWave is unique in its ability to segregate glitches from signals because the glitch/gaussian noise models provide an alternative to the signal model when the data contain instrument artifacts. Having a noise model that includes the possibility of glitches sets a higher bar for a candidate event to be classified as a signal. If the signal model leaves behind any significant excess power the glitch model will be preferred. While glitches which fortuitously end up being coincident in time and frequency pose a challenge for searches, BayesWave has successfully classified these transient events as being instrument noise. It should be pointed out that the glitch rejection fidelity will increase with the strength of the signal. High signal-to-noise ratio glitches place more stringent demands on the signal model because even small mismatches between the model and the data will leave behind significant residual while the glitch model, by construction, will account for all excess power. These results suggest that BayesWave can dramatically reduce the background for burst searches in Advanced LIGO/Virgo thereby increasing the range of the interferometers.

Glitch rejection is of no use if good candidates for detection are rejected as well. We confirm that BayesWave's detection capabilities are not hampered by the glitch model by analyzing data from LIGO's sixth science run to which we have added simulated gravitational wave signals of varying strength, to measure at what signal to noise ratio the injections are identified as signals by our algorithm. We find that BayesWave's detection efficiency behaves exactly as expected – when there is detectable power in multiple interferometers, the gravitational wave model is the preferred hypothesis for the data.

The Advanced LIGO detectors are due to be completed and collecting data before upgrades to Virgo are finished. The LIGO detectors are similarly aligned with one another and, as a consequence, are predominantly sensitive to a single polarization. Our assumption of elliptically polarized signals is most often sufficient for LIGO-only analyses and will therefore be adequate for the first advanced detector observations. However, for very high SNR injections of unpolarized signals ($\gtrsim 50$) the elliptically polarized signal model can not find an adequate fit and falsely classifies the injection as a glitch. While such loud signals are unlikely for near-future ground based detectors, we need to relax the requirement that signals be elliptically polarized without introducing unconstrained directions in the parameter space. If we were to allow h_+ and $h_\times$ to vary independently, most sky-locations would allow one polarization to be unconstrained which will negatively impact the Markov chain performance. Additional priors will be required on the signal model to hold both polarizations in check. A method to allow independent polarizations for the two-detector network without affecting MCMC convergence is under investigation, and BayesWave's glitch rejection capabilities will need to be re-quantified once we have relaxed this restriction on the signal model.

The results reported in this paper are suggestive that burst detection efficiency will benefit by employing BayesWave as a follow-up to background events and candidate detections from burst search pipelines such as cWB. This study alone is not sufficient to make that claim. A large-scale comparison between the detection efficiency of the burst pipeline with and without BayesWave's input is required to unambiguously quantify any improvement. Completing this comparison is of upmost importance prior to the start of advanced LIGO observations and is now under way.

The parameter estimation capabilities of BayesWave are in their infancy. Any and all quantitates of interest stem directly from the reconstructed waveform, and we have shown high fidelity between what BayesWave recovers and what is injected into the data. Now that a robust tool for assessing waveform characteristics is in place we have to develop a set of useful metrics for learning about the astrophysical nature of a burst source. These will most likely depend on different hypotheses for the GW emission. We have begun to explore which are the most useful waveform diagnostics by studying a wider variety of predicted burst waveform morphologies including core-collapse supernova, stellar mass black hole mergers, eccentric black binary encounters, etc.

V. ACKNOWLEDGMENTS

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-[DCC-XXXXXXXX]

[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] G. M. Harry (LIGO Scientific Collaboration), Class. Quantum Grav. **27**, 084003 (2010).
[4] *Advanced Virgo Baseline Design* (2009), document VIR-027A-09 in `https://tds.ego-gw.it/itf/tds/file.php?callFile=VIR-0027A-09.pdf`.
[5] N. Christensen and R. Meyer, Phys.Rev. **D64**, 022001 (2001), gr-qc/0102018.
[6] J. Veitch, V. Raymond, B. Farr, W. M. Farr, P. Graff, S. Vitale, B. Aylott, K. Blackburn, N. Christensen, M. Coughlin, et al., arXiv:1409.7215 (2014), 1409.7215.
[7] S. Klimenko, I. Yakushin, A. Mercer, and G. Mitselmakher, Class. Quantum Grav. **25**, 114029 (2008), arXiv:0802.3232.
[8] F. Cavalier, M. Barsuglia, M.-A. Bizouard, V. Brisson, A.-C. Clapson, et al., Phys.Rev. **D74**, 082004 (2006), gr-qc/0609118.
[9] S. Fairhurst, New J.Phys. **11**, 123006 (2009), 0908.2356.
[10] J. Markowitz, M. Zanolin, L. Cadonati, and E. Katsavounidis, Phys.Rev. **D78**, 122003 (2008), 0810.2264.
[11] R. Essick, S. Vitale, E. Katsavounidis, G. Vedovato, and S. Klimenko (2014), 1409.2435.
[12] T. Summerscales, A. Burrows, C. D. Ott, and L. S. Finn, Astrophys.J. **678**, 1142 (2008), 0704.2157.
[13] C. Rover, M.-A. Bizouard, N. Christensen, H. Dimmelmeier, I. S. Heng, et al., Phys.Rev. **D80**, 102004 (2009), 0909.1093.
[14] J. Logue, C. Ott, I. Heng, P. Kalmus, and J. Scargill, Phys.Rev. **D86**, 044023 (2012), 1202.3256.
[15] M. C. Edwards, R. Meyer, and N. Christensen, Inverse Prob. **30**, 114008 (2014), 1407.7549.
[16] B. Abbott et al. (LIGO Scientific Collaboration), Rept.Prog.Phys. **72**, 076901 (2009), 0711.3041.
[17] L. Blackburn, L. Cadonati, S. Caride, S. Caudill, S. Chatterji, et al., Class.Quant.Grav. **25**, 184004 (2008), 0804.0800.
[18] J. Aasi et al. (VIRGO Collaboration), Class.Quant.Grav. **29**, 155002 (2012), 1203.5613.
[19] N. J. Cornish and T. B. Littenberg, arXiv:1410.3835 (2014), 1410.3835.
[20] T. B. Littenberg and N. J. Cornish, arXiv:1410.3852 (2014), 1410.3852.
[21] D. Gamerman, *Markov Chain Monte Carlo: Stocastic Simulation of Bayesian Inference* (Chapman and Hall, London, 1997).
[22] M. Principe and I. M. Pinto, Class. Quantum Grav. **25**, 075013 (2008), arXiv:0806.4574.
[23] M. Principe and I. M. Pinto, Class. Quantum Grav. **26**, 045003 (2009).
[24] J. Crowder and N. Cornish, Phys. Rev. D **75**, 043008 (2007), arXiv:astro-ph/0611546.
[25] T. B. Littenberg and N. J. Cornish, Phys. Rev. D **80**, 063007 (2009), arXiv:0902.0368.
[26] T. B. Littenberg, Phys.Rev. **D84**, 063009 (2011), 1106.6355.
[27] P. Green, *Highly Structured Stochastic Systems* (Oxford University Press, 2003).
[28] P. Goggans and Y. Chi, *Bayesian Inference and Methods in Science and Engineering* (American Institute of Physics, USA, 2004).
[29] R. Swendsen and J. Wang, Phys. Rev. Lett. **57**, 2607 (1986).
[30] J. Abadie et al. (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. D **85**, 122007 (2012), 1202.2788.
[31] B. Abbott et al. (LIGO Scientific Collaboration), Phys.Rev. **D80**, 102001 (2009), 0905.0020.