

# Bayesian inference for spectral estimation of gravitational wave detector noise

Tyson B. Littenberg

*Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA)  
& Department of Physics and Astronomy, Northwestern University,  
2145 Sheridan Road, Evanston, Illinois 60208, USA*

Neil J. Cornish

*Department of Physics, Montana State University, Bozeman, Montana 59717, USA*  
(Received 21 October 2014; published 14 April 2015)

Gravitational wave data from ground-based detectors is dominated by instrument noise. Signals will be comparatively weak, and our understanding of the noise will influence detection confidence and signal characterization. Mismodeled noise can produce large systematic biases in both model selection and parameter estimation. Here we introduce a multicomponent, variable dimension, parametrized model to describe the Gaussian-noise power spectrum for data from ground-based gravitational wave interferometers. Called BayesLine, the algorithm models the noise power spectral density using cubic splines for smoothly varying broadband noise and Lorentzians for narrow-band line features in the spectrum. We describe the algorithm and demonstrate its performance on data from the fifth and sixth LIGO science runs. Once fully integrated into LIGO/Virgo data analysis software, BayesLine will produce accurate spectral estimation and provide a means for marginalizing inferences drawn from the data over all plausible noise spectra.

DOI: [10.1103/PhysRevD.91.084034](https://doi.org/10.1103/PhysRevD.91.084034)

PACS numbers: 04.30.-w, 04.80.Nn, 95.55.Ym

## I. INTRODUCTION

Gravitational wave data from ground based observatories such as LIGO [1] and Virgo [2] are dominated by instrument noise. Even with the improved hardware capabilities of the Advanced LIGO/Virgo interferometers [3,4] sensitive analysis methods must be in place to take full advantage of the data. While a wide variety of search pipelines exist for a similarly wide variety of GW source types (e.g. [5–7]), all fundamentally stem from the same principle—assessing whether the detector output is statistically different from our model for the data. For an examination of the inherent similarity of different GW search techniques see Ref. [8].

For analyses which have the advantage of strong theoretical predictions for the gravitational wave signal, such as searches for the inspiral and merger of stellar mass compact binaries (neutron stars and/or black holes) template based techniques are optimal [9,10]. In the template based analyses used for signal characterization [11,12], a trial waveform, or “template,” is subtracted from the data and the residual is compared to a theoretical model for the instrument noise. The statistic used to assess whether the residual and the noise model are consistent is the likelihood, or the probability density of the detectors producing data  $d$  for a hypothetical gravitational wave signal  $h$  with parameters  $\theta$ . Assuming we understand the properties of the noise distribution, the likelihood will be maximized when the template waveform matches the true signal

modulo statistical fluctuations due to the specific realization of the noise.

Any unsound approximations in the likelihood will contribute to systematic errors which could result in erroneous inferences as the template waveform parameters flex away from the “true” values in an attempt to achieve the expected statistics for the residual. Receiving most of the spotlight in the effort to control systematic errors are the template waveforms themselves [13–17]. However, understanding the noise is just as important if we want to ensure accurate astrophysical parameter estimation and discoveries, or limits pertaining to fundamental physics such as alternative theories of gravity [18–20] and the equation of state of neutron stars [21–24].

The LIGO/Virgo noise is dominated by three main components: Seismic noise steeply limits sensitivity below 10 Hz; thermal noise from the mirror suspensions and coatings dominate between 10 and 200 Hz; and quantum (photon shot) noise is the limiting factor at high frequencies. Each component is a broadband effect with relatively simple frequency dependence. Standing out above the broadband noise are high power, 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 data calibration.

Approaches for improving the noise model include modifying the functional form of the likelihood [25–28]; demanding additional consistency checks between analyses on single detectors and the network [29,30]; or introducing

degrees of freedom which help fit to the instrument noise [31,32]. Here we introduce a new way of parametrizing LIGO/Virgo Gaussian noise and use modern data analysis methods for fitting to, and eventually marginalizing over, the noise model. A companion publication will address the nonstationary and non-Gaussian noise by introducing the BayesWave algorithm [33]; a trans-dimensional Bayesian approach similar in spirit to BayesLine but employing a different basis—wavelets—to model transient, non-Gaussian noise. While demonstrations in this paper use BayesLine as a “preprocessor” spectral estimation code, the algorithm is too computationally intensive to be employed in low latency searches for GWs. BayesLine will ideally be integrated into the parameter estimation pipelines used to followup candidate transient GW events. Doing so will allow estimated parameter uncertainties to incorporate statistical errors due to uncertainty in the noise power spectral density. It may also prove beneficial to adapt BayesLine for stochastic GW background searches.

In Sec. II we motivate the need for real-time spectral estimation, in Sec. III we consider features commonly found in LIGO and Virgo noise spectra which place different demands on our model. Section IV introduces the BayesLine algorithm and we demonstrate its performance on LIGO data in V. Our conclusions are summarized in Sec. VI while Appendix A briefly explores the prospects of coherently fitting out spectral lines associated with the electrical power supply.

## II. MOTIVATION

The noise from gravitational wave detectors  $\mathbf{n}$  is approximately Gaussian with zero mean (i.e.  $\langle \mathbf{n} \rangle = 0$ ) and stationary (meaning the noise statistics are not time dependent). Here, and throughout, bold-faced quantities indicate the vector of all samples, while unbolded quantities refer to a single sample. Because the autocorrelation matrix of stationary noise depends only on the lag, the noise covariance matrix of  $\tilde{\mathbf{n}}$ , where the tilde represents a Fourier transform, is diagonal (i.e., there are no correlations between noise at different frequencies). Under these conditions we can completely characterize the noise by its one-sided power spectral density (PSD)  $S_n(f) \equiv \frac{2}{T} \langle |\tilde{\mathbf{n}}(f)|^2 \rangle$  where  $T$  is the duration of data being analyzed, typically 10 s to 1000 s for LIGO/Virgo compact binary signals. In the paradigm of stationary and Gaussian noise we arrive at the standard likelihood used in gravitational wave data analysis [34]

$$\log p(\mathbf{d}|\boldsymbol{\theta}) = - \sum_I^{\text{Net}} \left[ \int_0^\infty \frac{|\tilde{\mathbf{r}}^I(f; \boldsymbol{\theta})|^2}{S_n^I(f)} df \right] + C \quad (1)$$

where  $\tilde{\mathbf{r}}(f; \boldsymbol{\theta}) = \tilde{\mathbf{d}}(f) - \tilde{\mathbf{h}}(f; \boldsymbol{\theta})$  is the residual noise after the gravitational wave template has been regressed from the

data, and the summation is over detectors in the network (e.g., the two LIGO detectors and Virgo).

Characterizing the detector noise is a challenge because it is not stationary for times much longer than a few tens of second, depriving us of a sufficiently accurate “reference” noise spectrum that can be used for all analyses. Instead, the noise floor must be estimated from the data directly. Different analyses have different prescriptions for PSD estimation. Past searches for binary in-spiral signals used a running average of the instrument power spectrum over long durations of time [6] while the parameter estimation follow-up analysis relied on Welch averaging the power spectrum of many segments of data near in time to a candidate signal coming from the search pipeline [11,12]. Searches for unmodeled signals have their own methods for PSD estimation, including pre-conditioning the data by filtering narrow-band spectral features and piece-wise whitening the Fourier domain data [35], and whitening across frequency layers in a discrete wavelet transform representation of the data [36].

Consider the parameter estimation follow-up pipeline approach to PSD estimation. Relying on off-source data to estimate the noise level presents several hazards including times when the detector noise experiences large excursions due to impulsive noise events, “glitches,” or long duration fluctuations of data quality due to prolonged changes in the environment such as seismic activity, wind, etc. In past analyses, times used to produce PSDs were manually inspected to ensure quality data was employed in the Welch averaging and the windows of time were adjusted accordingly. Finding adequate data for noise estimation could be challenging, and the posterior distribution functions recovered occasionally showed statistically significant dependence to the data used for PSD estimation [11,32].

The challenge of producing good initial estimates of the noise power spectrum will be compounded in the advanced detector era. As the sensitivity of the detectors improves, particularly at low frequency, the amount of time that a gravitational wave chirp signal will be in band increases. For example, a 1.4–1.4  $M_\odot$  binary neutron star merger takes  $\sim 25$  s to evolve from 40 Hz (the approximate low frequency cutoff during the initial science runs) to the point where the stars merge. Advanced LIGO/Virgo are expected to ultimately reach down to frequencies of  $\sim 10$  Hz [37], at which point the binary is in band for  $\sim 10^3$  s. Reliance on averaging for PSD estimation will demand  $\sim 10^4$  s—h—of stationary noise similar to the data containing the candidate event. The demands on detector stability for the durations needed to achieve good noise estimates are impractical.

It was demonstrated in Ref. [32] that sensitivity to the initial PSD estimation could be mitigated, but not eradicated, by marginalizing over the overall power in the instrument background using parameters which produced a step-wise rescaling of the noise spectrum. The PSD rescaling is a step

in the right direction, and is computationally inexpensive, but requires an initial estimate of the PSD and is limited in its flexibility; it was unable to modify the spectral slope or the location and relative amplitude of the spectral lines found in LIGO/Virgo data.

Given this reality, an alternative approach to determining the PSD—one which only relies on the data containing the trigger—is desirable. Ideally we would have a parametrized model for the instrument noise, just as we do for the gravitational wave signal, and the two will be deduced simultaneously during the parameter estimation analysis. Bayesian spectral estimation and parameter estimation is a lively topic of research. We refer the interested readers to Ref. [38] for a thorough primer on the subject. Here we present a novel approach to PSD estimation where we introduce a variable number of phenomenological parameters which are sufficiently flexible to fully characterize the noise power spectrum but are held in check using Bayesian model selection to prevent over-fitting the data. The PSD model is broken up into two components: Cubic spline interpolation is used to fit to the broadband spectrum dominated by seismic, thermal, and shot noise while a linear combination of Lorentzians is used to whiten the narrow-band resonant features in the data due in part to the AC power supply, calibration lines, and vibrational modes of the suspension system supporting the test masses.

### III. THE LIGO/VIRGO POWER SPECTRUM

Before launching into the details of how we will model the PSD, and how this improvement impacts inferences drawn from the data, we will first motivate the need for more sophisticated treatment of the instrument noise by demonstrating the breakdown of the standard assumptions about the data being stationary on timescales needed to perform in-spiral analyses. To do so we will use data from the LIGO Livingston Observatory to illustrate the challenges facing future observations. To start, we will begin with an example Fourier power spectrum of typical LIGO data shown in the gray (dashed) line of Fig. 1. Plotted over the spectrum is the design sensitivity (blue, dashed line) which is the type of PSD typically used in studies relying on simulated data. Notice that the design sensitivity curve is missing many important features found in real data, particularly the high-power, narrow-band features, or “lines,” located throughout the sensitivity band of the detector. Finally, the red (solid) curve shows a typical PSD that would be used during a parameter estimation follow-up analysis of a candidate chirp signal. To demonstrate the challenges of estimating the PSD using the methods employed in deriving the red curve, we will dissect the LIGO power spectrum into two components: The broadband noise as approximated by the design sensitivity curves, and the spectral lines found in the real data.

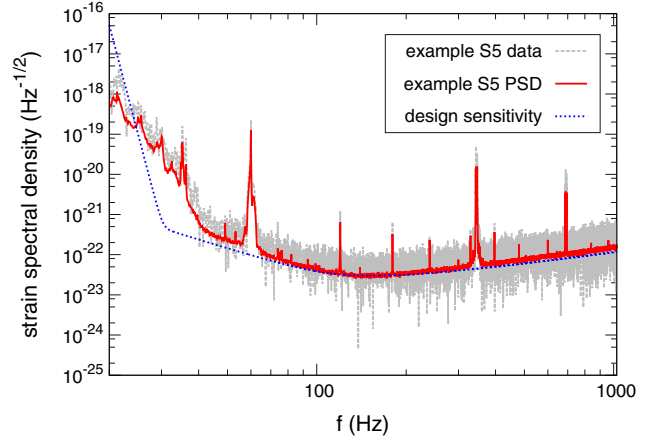


FIG. 1 (color online). Example strain spectral density (square root of the Fourier power; gray, dashed) and  $\sqrt{2S_n(f)/T}$  (red, solid) during LIGO’s 5th science run, compared to the design sensitivity (blue, dotted). Notice that the design sensitive consists strictly of broadband features due to seismic, thermal, and quantum noise, while the actual data show high-amplitude narrow-band spectral lines. The design sensitivity is the type of noise curve used for studies in simulated data, while the red (solid) line is needed for analyzing the actual data.

#### A. The broadband LIGO/Virgo noise

Ground based gravitational wave interferometer noise is dominated by three main components (details of the bandwidth for each effect depend on the instrument): Seismic noise steeply limits sensitivity below a few tens of Hz, thermal noise from the mirror suspensions and coatings dominate between the seismic wall and a few hundred Hz, and quantum (photon shot) noise is the limiting factor at high frequencies. Each component is a broadband effect with relatively simple behavior as a function of frequency, the sum of which produce the familiar sensitivity curves (e.g. the blue dotted line in Fig. 1) where the three different slopes are where the three different components—seismic ( $f \lesssim 30$  Hz), thermal ( $30 \lesssim f \lesssim 100$  Hz), and shot noise ( $f \gtrsim 100$  Hz)—dominate the noise budget. The example shown in Fig. 1 comes from 1024 s of data taken toward the end of LIGO’s fifth science run—when LIGO achieved initial design sensitivity—which took place from November 2005 to October 2007, henceforth “S5.”

Upgrades to the LIGO/Virgo facilities are, in a broad sense, intended to further suppress each of these three components. A detailed discussions of the LIGO noise budget during S5 is found in Refs. [39,40]. For a preview of the expected noise budget for Advanced LIGO see Fig. 2 of Ref. [37]. Between the completion of S5, and before the full-scale commissioning work for Advanced LIGO began, the S6 science run took place between July 2009 to October 2010. S6 achieved better sensitivity than S5 but the power spectrum was qualitatively similar to Fig. 1.

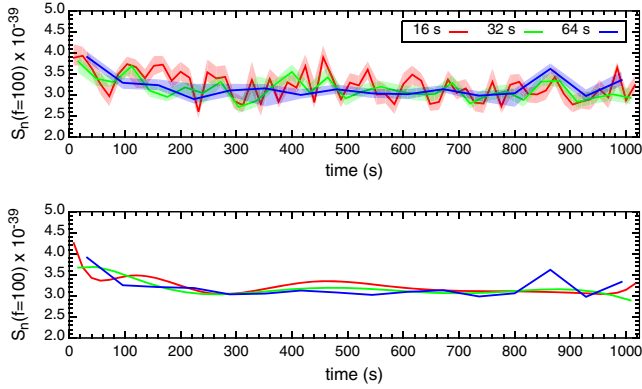


FIG. 2 (color online). Demonstration of the time variability for the broadband Gaussian noise. Top panel: PSD over 1024 s of S6 data at 100 Hz estimated using BayesLine over intervals of 16 (red), 32 (green), and 64 (blue) seconds of data. The thick line is the median PSD from each segment’s chain and the shaded region spans the 90% credible interval. Bottom panel: Legendre polynomial fit to the data, where a RJMCMC was used to determine the best-fit degree ( $n$ ) for the polynomial ( $n = 10$  for 16 s,  $n = 7$  for 32 s, and  $n = 11$  for 64 s).

For the remainder of this article, all examples will use data taken during S6.

To quantify how nonstationary the noise can be we begin with 1024 s of data which is the nominal duration used for PSD estimation in the follow-up of S6 triggers [11,12]. We then divide that data into equal duration segments of width 32 s—the amount of data which would be needed to perform parameter estimation of a binary neutron star signal in S6. In each 32 s segment of data we use the BayesLine algorithm (a detailed description of which can be found in Sec. IV) to compute  $S_n(f)$  at 100 Hz, with  $1\sigma$  error bars as characterized by a Markov chain. The PSD at 100 Hz is then plotted for each segment, as a function of segment start time relative to the full 1024 s of data, as the gray (dashed) points in Fig. 2. We then repeat the analysis using 16 s (red, solid) and 64 s (blue dotted) seconds of data. Notice that over the span of 1024 s desired amount of data used to compute the average PSD during S6—we find statistically significant variation in the PSD. To quantify the time dependence, we fit the time-dependent PSD to an  $n$ th-order polynomial. The curves in Fig. 4 show the best-fit polynomials whose order  $n$  (10th for the 16 s segments, 7th for the 32 s segments, and 11th for the 64 s segments) had the highest Bayesian evidence as determined by a Reverse Jump Markov chain Monte Carlo (RJMCMC) code. During the advanced detector era, a binary neutron star signal will be in band for this entire 1024 s segment of data and, using the same procedure as in previous analyses, nine hours of data would be required to estimate the PSD. As can be seen from this demonstration, approximating the noise as being stationary over such long-duration intervals is not reliable. By contrast, BayesLine only needs the data being analyzed for a GW signal to determine the PSD, as opposed to

requiring durations of data many times longer for spectral estimation.

## B. The spectral lines

Above the broadband noise floor for ground-based interferometric detectors are comparatively high power, narrow band, spectral features as seen in Fig. 1. The most significant spectral lines in Initial LIGO/Virgo data originate from one of three main processes: Resonances of the wires which suspend the mirrors in the interferometer, henceforth referred to as “violin” modes, predominantly located around 320 Hz (plus higher harmonics) in initial LIGO; the AC electrical supply “power line” and subsequent higher frequency harmonics (at 60 Hz in the LIGO observatories); and calibration lines which are “injected” into the data by driving the mirrors at specific frequency and amplitude. Advanced LIGO/Virgo will have similar features, though the central frequencies may differ. The precise frequency and amplitude of the environmental (power and violin) lines tend to drift on timescales shorter than the amount of data needed to estimate PSDs for in-spiral signals through averaging off-source data.

Figure 3 depicts the time-evolution of the frequency (left axis, red solid line) and amplitude (right axis, blue dashed line) of the 60 Hz line computed by taking 1024 s of data and dividing it into 128 equal-length (8 s) segments. Each segment of data was then Fourier transformed, and the frequency and amplitude of the peak in the power spectrum found in a window around 60 Hz was recorded. Higher harmonics (e.g. at 120 Hz) show near-perfect correlation with the 60 Hz power line.

As is evident in Fig. 3, the amount of fluctuation in the location of spectral lines can be dramatic—migrating by tens of Fourier bins for inspiral analyses—which results in




FIG. 3 (color online). The 60 Hz power line shows nonstationary frequency (left axis, red solid line) and amplitude (right axis, blue dashed line) over timescales shorter than the amount of data needed to characterize binary in spiral signals. Higher frequency harmonics are practically 100% correlated with the main power line.

differences between the power spectrum used for the noise model and the actual power spectrum of the data.

In principal, one could argue that notch-filtering the lines [i.e., excluding frequencies associated with line features from the likelihood integration in Eq. (1)] would be a simpler strategy. However, what is not captured in Fig. 3 are similar fluctuations in the line width. Any notch filter would need to be conservative in order to ensure that no line power would leak into the integration domain and doing so would be at the expense of potentially throwing away good data. A “real time” strategy for mitigating spectral lines is therefore preferred.

Several methods for dynamically removing lines have been put forward. Analyses by Finn and collaborators using the 40 m prototype LIGO interferometer explored the use of Kalman filters [41,42] to subtract the violin modes from the data, while relying on auxiliary magnetometer data to approximate the contribution from the power lines in the gravitational wave channel and subsequently subtract the lines from the data. Sintes and Schutz [43] used data from the Glasgow laser interferometer to develop the *coherent line removal* algorithm which detects (and regresses) lines by searching for coherent harmonics of narrow-band features.

We will now introduce an alternative method for inferring  $S_n(f)$  which is generically applicable to different spectral features in the data and does not depend on off-source data or long-term averages. The algorithm is constructed with the end goal of simultaneously characterizing gravitational wave signals and detector noise in mind. To do so we have created a framework by which parameters for the Gaussian component of LIGO/Virgo noise can be seamlessly integrated into the existing parameter estimation software. The algorithm, named BayesLine, uses RJMCMC methods to determine the most probable model for the PSD. What follows is a detailed description of the algorithm and examples designed to demonstrate its benefits for LIGO/Virgo data analysis.

#### IV. THE BAYESLINE ALGORITHM

To infer the noise power spectral density from the data we have constructed a parametrized model for  $S_n(f)$  which uses a MCMC to determine the posterior distribution function for the PSD. MCMC methods have become increasingly common in astrophysics and are ubiquitous within the gravitational wave community, so in lieu of rehashing the usual description of the algorithm we refer the reader to an accessible subset of the available literature [44–47].

Just as our demonstrations of nonstationary LIGO noise separated the broadband noise from the lines, we will model the two components separately, so that the total PSD  $S_n(f)$  is the sum of the broad band “smooth” part of the power spectrum and the line model. The dimension of the broadband noise  $S_S(f; \xi, N_S)$  and the spectral line models

$S_L(f; \lambda, N_L)$ —with parameters  $\xi$ ,  $N_S$ ,  $\lambda$ , and  $N_L$  to be defined in the next subsections—is a free parameter, with the quantities  $N_S$  (the number of control points in the cubic spline) and  $N_L$  (the number of Lorentzians) tracking each component’s dimension.

For this application,  $N_S$  and  $N_L$  are nuisance parameters so we are not interested in selecting the PSD model with the largest evidence. Instead, because we use an RJMCMC to sample the entire model space, we produce the model averaged distribution for  $S_n(f)$ . As a result, inferences drawn from analyses which use BayesLine will be marginalized over the noise model.

##### A. The cubic spline model for broadband noise

The smooth part of the PSD,  $S_S(f; \xi, N_S)$ , is fit using cubic spline interpolation. The parameters of the model are the number  $N_S$ , and location in frequency-PSD plane  $\xi_i \rightarrow \{f_j, S_j\}$  of each spline point  $i \in [0, N_S]$  in the interpolation. A trial set of control points  $\{\xi_j\}$  is chosen randomly from the prior and fed into the cubic spline routines available in the *GNU Science Library* [48] to determine the spline coefficients  $c_n$ . From the coefficients we compute the cubic spline curve that goes through all of the points

$$S_S(f_j \leq f < f_{j+1}) = \sum_{n=0}^3 c_n^j (f - f_j)^n \quad \text{for } j \in [0, N_S]. \quad (2)$$

That curve will then be compared to the Fourier power spectrum of the actual data when we compute the likelihood for the trial set of spline points as described in Sec. IV C.

The spline control points start on a regular grid in frequency with an interval of  $\sim 10$  Hz and  $S_j$  is initialized with the median Fourier power in a small window around frequency  $f_j$ . During the Markov chain analysis the location of each spline point is adjusted both in frequency and power spectral density, existing points can be removed from the fit, and new points can be added.

Priors for the spline model are uniform in frequency and in power spectral density. The maximum number of control points is set by the total bandwidth of the data and the initial  $\sim 10$  Hz grid. The range of allowed PSD values for each spline point depends on the type of data being analyzed, as we envision BayesLine being a useful tool not only for gravitational wave data but also for auxiliary channels used to monitor the external environment at each observatory.

##### B. The Lorentzian line model

The most noticeable lines in the power spectrum are the “violin modes” caused by resonances in the mirror suspension system. A good model for the source of the suspension lines is a noise-driven and damped harmonic oscillators. The power spectrum for such systems is known to be a Lorentzian [41] and so we will choose that as the

basis for our model for the lines. While the physical mechanism responsible for other spectral features found in the data—including the 60 Hz power line (plus harmonics) and calibration lines—are not best represented by Lorentzians, we found that our choice of basis was adequate for all spectral features in the data and so opted for the simplest model to implement. A more sophisticated model which is tailored for the different line-inducing processes in the data is discussed in the Appendix and left as a future direction to improve the BayesLine algorithm.

The details of our Lorentzian line model are as follows: We begin with a parameter vector  $\lambda_i \rightarrow \{A_i, f_i, Q_i\}$  for the  $i^{\text{th}}$  Lorentzian in the fit where  $A_i$  is the amplitude,  $f_i$  is the central frequency, and  $Q_i$  is the quality factor which governs the width of the line.

The full line model  $S_L(f)$  is then the sum of each individual Lorentzian  $\Lambda(f; \lambda_i)$  in the fit

$$S_L(f) = \sum_i^{N_L} \Lambda(f; \lambda_i)$$

$$\Lambda(f; \lambda_i) = \frac{z(f)A_i f_i^4}{(f_i f)^2 + Q^2(f_i^2 - f^2)^2} \quad (3)$$

where  $z(f)$  is used to truncate the tails of the Lorentzian distribution via

$$z(f) = \begin{cases} 1, & \text{if } |f - f_i| \leq \delta f \\ e^{-\frac{|f - f_i| - \delta f}{\delta f}}, & \text{otherwise} \end{cases} \quad (4)$$

with  $\delta f \sim f_i/50$  used to truncate the lines after several intervals of the full width at half maximum bandwidth for the line. The precise choice of  $\delta f$  was based on tests using LIGO data and a wide range of definitions for  $\delta f$  worked equally well.

Our priors for the central frequency of the spectral lines are uniform over the full bandwidth of the data and deserve further improvements: Archived and auxiliary data provide information about the location of spectral lines; the power line (and its harmonics) are within a fraction of a percent of 60 Hz (and higher multiples); calibration lines are purposefully added to the data at known frequency and amplitude; and the suspension lines are mechanical resonances with narrowly constrained frequencies. The “signal to noise” of the spectral lines is very high and will overwhelm our choice of prior, but making more informed selections about the spectral line parameters will improve the convergence of the BayesLine model when integrated with algorithms which rely on parallel tempering for evidence calculations [12,33].

### C. The likelihood

We arrive at the full noise power spectral density by adding the two components together

$$S_n(f) = S_L(f; \lambda, N_L) + S_S(f; \xi, N_S). \quad (5)$$

To construct the likelihood we consider the joint probability that the  $N$  complex Fourier coefficients of the data  $\mathbf{d}$  would be realized from  $S_n(f)$  assuming that the whitened data  $\tilde{\mathbf{d}}(f)/(\frac{T}{2}S_n(f))^{1/2}$  should be consistent with a unit normal distribution

$$\log p(\mathbf{d}|\lambda, N_L, \xi, N_S) = -\frac{2}{T} \sum_f^N \left[ \frac{|\tilde{\mathbf{d}}(f)|^2}{S_n(f)} \right] \quad (6)$$

where we have left off the constant normalization  $-\frac{N}{2} \log(2\pi)$  as the MCMC is only concerned with relative likelihoods between points in parameter space.

### D. Reverse Jump Markov Chain Monte Carlo

The central engine of BayesLine is a trans-dimensional variant of a Markov chain—RJCMC—where the chain transitions between different models for the data, potentially of different dimension, thereby simultaneously characterizing the posterior distribution function for each model and producing the evidence ratio, or Bayes factor, for each model [49]. RJCMC methods have been used in gravitational wave astronomy to assess whether simulated data contain a detectable astrophysical signal [31,50–54], to estimate the mass distribution of stellar-mass black holes [55], and as a proof of principle for modeling non-Gaussian noise [31].

RJCMC has the unique ability to transition between competing models, effectively making the model one of the search parameters. Like its fixed dimension counterpart, the RJCMC is guaranteed to (eventually) converge to the true target distribution—the likelihood distribution across model space. The marginalized likelihood, or relative evidence, for each model is the number of iterations the chain spends in each model divided by the total number of iterations in the chain. A derivation of the error in RJCMC evidence calculations is in our companion BayesWave paper, Ref. [33].

Allowing for the exploration of different models (which may differ in dimension) requires a separate Metropolis-Hastings step which proposes to move the chain from one model to another. Parameters  $\theta_i$  for trial model  $\mathcal{M}_i$  are drawn from  $q(\theta_i|\mathcal{M}_i)$ . If the models are nested, such as proposing additional parameters to include in the existing set, all of the like parameters are held fixed while the new parameters are drawn from  $q$ . Once the new models parameters are in hand the trans-model Hastings ratio is calculated by

$$\mathcal{H}_{\mathcal{M}_i \rightarrow \mathcal{M}_j} = \frac{p(\mathbf{d}|\theta_j, \mathcal{M}_j)p(\theta_j|\mathcal{M}_j)q(\theta_i|\mathcal{M}_i)}{p(\mathbf{d}|\theta_i, \mathcal{M}_i)p(\theta_i|\mathcal{M}_i)q(\theta_j|\mathcal{M}_j)} |J_{ij}| \quad (7)$$

where the Jacobian  $|J_{ij}|$  accounts for any change in dimension between models  $\mathcal{M}_i$  and  $\mathcal{M}_j$ . Selecting an

efficient proposal distribution for model transitions is typically the major obstacle in the implementation of an efficient RJ routine. If the proposal distributions yield the model parameters directly, instead of a set of random numbers which are then used to determine the new model parameters, the Jacobian is unity and can be neglected.

For the spline model uniform draws from the prior for the frequency and power spectral density of control points provide adequate mixing of trans-dimensional proposals, though this would improve with stricter priors. For the line model we again rely on uniform draws from the prior for the amplitude and quality factor for the Lorentzians. For the central frequency the power spectrum is first divided into several narrow-band segments which are weighted by the power in that segment divided by the median power of the full data, thereby preferring to propose Lorentzians in segments that have a large excess of power. Segments are proposed according to their relative weighting and the frequency is drawn uniformly from within the narrow-band segment. We verify that detailed balance is satisfied by running BayesLine with the likelihood set to a constant and confirming that the recovered posterior agrees with our priors.

**V. STUDIES USING LIGO DATA FROM S6**

With the likelihood and prior distributions defined, we can now completely characterize the posterior distribution function of the PSD using the RJMCMC algorithm. We will now make use of this formalism to demonstrate how our parametrized PSD outperforms the off-source averaging method for spectral estimation using data acquired by the LIGO observatories during S6.

For the stopping condition of the Markov chain we determine when the chain has converged by using a posterior predictive check which demands that the whitened data (which should be consistent with a normal distribution)




FIG. 4 (color online). Example S6 data power spectrum (gray, dashed) with BayesLine PSD shown separated into the cubic spline (red, solid) and Lorentzian (blue, dotted) components.

do not have large-sigma outliers. We do not currently monitor the autocorrelation length to adaptively determine when enough posterior samples have been collected after the convergence criterion is satisfied, instead using a fixed number of samples in the chain post-convergence that has proven to work well. As BayesLine is incorporated into existing LIGO/Virgo Bayesian inference pipelines the stopping criteria will need to be more carefully handled.

Run-times for the current BayesLine implementation are a factor of a few times larger than the duration of data being analyzed, depending on the sampling rate. For example,




FIG. 5 (color online). Initial LIGO example distribution of whitened Fourier domain data using the standard PSD estimation in LALInference (blue, dotted) and the BayesLine PSD (red, solid), compared to a normal distribution with zero mean and unit variance (gray, dashed) which is assumed by the likelihood function.




FIG. 6 (color online). Example distribution of whitened Fourier domain data made from the amount of data consistent with what is needed for Advanced LIGO BNS analyses. PSDs were determined via the standard PSD estimation in LALInference using 10 (magenta, dotted) and 30 (blue, short dashed) 1024 s segments of data for averaging and the BayesLine PSD (red, solid). All three distributions are compared to  $N[0, 1]$  (gray, long dashed).

BayesLine requires  $\sim 1$  hour to estimate the PSD for 1024 seconds of data sampled at 4096 Hz. The run time is insignificant when compared to the computational cost of compact binary parameter estimation analyses [12].

To begin we show in Fig. 4 a typical S6 LIGO power spectrum (gray, dashed) and the two component BayesLine maximum a posteriori PSD (with the appropriate  $T/2$  normalization). The spline model is shown in the red (solid) curve while the Lorentzians used to fit the spectral lines are depicted by the blue (dotted) lines. The full PSD would be sum of the spline and Lorentzian fit. For this example we use  $T = 32$  s of data from the LIGO Livingston Observatory. Our choice for  $T$  is consistent with the amount of data required for parameter estimation follow up of a binary neutron star (BNS) signal. BNS inspirals pose the most stringent challenge for PSD estimation (among the transient sources) because the duration of

the signal increases with decreasing mass. Binary neutron stars are anticipated to be the lowest mass signals and, as a result, require the longest data segments.

Using averaged off-source data to estimate the PSD would demand  $\gtrsim 1024$  s of data. To show the difference between the BayesLine noise spectrum and a PSD estimated by averaging we use the 1024 s immediately after the 32 s segment of data used in Fig. 4 to compute  $S_n(f)$  as was done during S6. Had the data from Fig. 4 contained a candidate detection, this would be indicative of the PSD used in the analysis. To compare the two PSDs we use each of them to whiten the original 32 s of data, and then histogram the real and imaginary Fourier coefficients.

According to the likelihood function used for parameter estimation [(Eq. (1))] the whitened data should be drawn from a normal distribution with zero mean and unit variance. Figure 5 shows the distribution of the whitened




FIG. 7 (color online). Comparison of marginalized posterior distribution functions of the mass ratio  $q = m_2/m_1$ , chirp mass  $M_{\text{chirp}} = (m_1 m_2)^{3/5} (m_1 + m_2)^{-1/5}$ , and luminosity distance  $D_L$  for two different simulated compact binary signals in 16 s of real detector noise using LALInference. The x-axes are shifted by the true value of each parameter. Results obtained with the BayesLine PSD (red, solid) are compared to posteriors using off-source PSDs determined from 256 s of data before the signal entered the LIGO band (gray, solid), 256 s after the merger time (dark blue, dashed) and 512 s after the merger time (blue, dotted). When the noise spectrum is stable the BayesLine PSD agrees well with off-source estimates (bottom row) and does not “harm” parameter estimation. Posteriors in the top row are sensitive to the start time of the data used for spectral estimation. The BayesLine posterior is most consistent with the results using data immediately after the trigger time. This result suggests that BayesLine will enable parameter estimates that are more robust to long-term nonstationary noise.



data using the BayesLine PSD (red, solid) and the off-source PSD (blue, dotted) as compared to a zero-mean unit-variance normal distribution  $N[0, 1]$  (gray dashed). The difference is striking—using the off-source PSD leaves behind significant large  $\sigma$  tails. Because the likelihood is maximized when the whitened residual looks Gaussian, the excess tails left behind by the off-source whitening invite bias by the gravitational wave model in an attempt to account for the non-Gaussian residual. On the other hand, our BayesLine PSD produces a Gaussian residual thereby suppressing the potential for large systematic errors in GW parameter estimation due to the PSD estimation.

As motivated in Sec. I, estimating  $S_n(f)$  for long-duration signals in the advanced detector era using off source data is not feasible. To illustrate this reality we repeat the analysis from Fig. 5 but using observation times  $T = 1024$  s consistent with a BNS signal in Advanced LIGO/Virgo. We estimate the PSD and histogram the whitened Fourier coefficients, shown in Fig. 6, using BayesLine (red, solid)  $10 \times 1024$  (magenta, small dots) and  $30 \times 1024$  (blue, large dots) seconds of off-source data. The 10240 s we consider the bare minimum for PSD estimation, while the 30720 s segment uses the same number of segments (30) in the Welch average as the S6 example for advance LIGO-like segment lengths. Note that  $\sim 30000$  s ( $\sim 8$  h) of stationary noise is unlikely to be available due to environmental disturbances affecting data quality.

To demonstrate how the PSD model can impact parameter estimation we add simulated binary black hole signals (with total mass  $\sim 15 M_\odot$ ) to 16 s of real data and use the LIGO/Virgo parameter estimation software LALInference to recover the signals and characterize their posterior distribution functions [12].

Figure 7 shows the poster distributions for the mass ratio  $q = m_2/m_1$ , chirp mass  $M_{\text{chirp}} = (m_1 m_2)^{3/5} (m_1 + m_2)^{-1/5}$ , and luminosity distance  $D_L$  for two different simulated compact binary signals in 16 s of real detector noise using LALInference. The x-axes are shifted by the true value of each parameters. The red solid distribution was obtained using the BayesLine PSD. It is plotted with posteriors inferred from the same data, of the same simulated source, but using off-source PSDs determined from 256 s of data starting (i) before the signal entered the LIGO band (gray, solid), (ii) 256 s (dark blue, dashed) and (iii) 512 s (blue, dotted) after the merger time. Different rows correspond to different signals added to different 16 s segments of data. When the noise spectrum is stable the BayesLine PSD agrees well with off-source estimates (bottom row) and does not “harm” parameter estimation. Posteriors in the top row are sensitive to the start time of the data used for spectral estimation. In each example, the BayesLine posterior is most consistent with the results using data immediately before (top row) or after (bottom row) the trigger time. This result suggests that BayesLine will enable

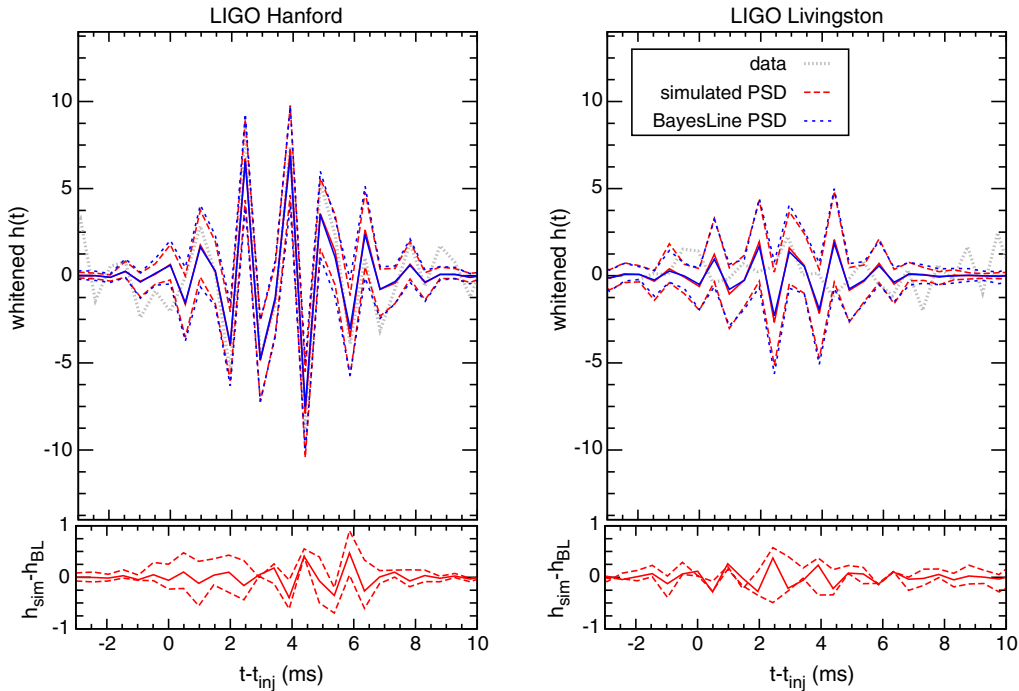


FIG. 8 (color online). Comparison of reconstructed waveforms in simulated data using the true PSD (red) and the BayesLine PSD (blue). The top panels are the whitened data (grey) and waveforms. Solid lines indicate the median reconstructed signal while dashed lines span the  $1\sigma$  errors. The bottom panel shows the whitened residual between the two recovered waveforms, which is consistent with zero. The figure demonstrates how including the BayesLine PSD parameters does not degrade signal recovery.

parameter estimates that are more robust to long-term nonstationary noise.

In principle there is a risk that the PSD model is able to fit-out part of the gravitational wave signal and thus introduces, instead of resolves, biases. In practice this is not a problem. Our ultimate vision of having BayesLine and gravitational wave parameters simultaneously in the model will not suffer from this hazard because the template provides a *much* better model for the GW signal than BayesLine and so the RJMCMC will, by construction, prefer models which do not interfere with the signal characterization. We demonstrate this point using simulated data (so the true PSD is known) and a simple model for the gravitational wave signal and template (a sine Gaussian). The data is analyzed once using the LIGO design sensitivity curve (which was used to simulate the noise), and again using BayesLine and the GW model simultaneously. We show the agreement between our fit and the true PSD, as well as between the recovered waveforms from the two runs in Fig. 8. The top panels show the whitened simulated data in gray, and the recovered waveforms are plotted using the true PSD (red) and the BayesLine PSD (blue). The solid lines are the median waveform and the dotted lines span the  $1\sigma$  errors for the reconstructed waveform. The agreement between the recovered waveforms using the different PSDs is excellent. The bottom panel shows the median and  $1\sigma$  residuals between the two reconstructed waveforms, which are consistent with zero.

## VI. CONCLUSIONS AND FUTURE GOALS

Analysis of ground-based gravitational wave interferometer data is challenging in part because the instrument noise, over intervals of time longer than duration of transient signals, is not stationary making spectral estimation a challenge. The standard approach to noise characterization used by analyses for binary in-spiral signals involves averaging over long durations of data compared to the amount of time that the signal is in band and is thereby susceptible to systematic errors. In this paper we provide an alternative approach to PSD estimation—dubbed BayesLine—for inferring the noise spectrum directly from the data being analyzed. Doing so within the MCMC framework means that our inferences about GW detections can be marginalized over all PSDs consistent with the data. The BayesLine algorithm has the added benefit of freeing the analyses from needing to use large amounts of data for spectral estimation.

Our parameter estimation results shown in Fig. 7 use a relatively small amount of data compared to the duration of signals expected during the Advanced LIGO era and still suggest that BayesLine will enable robust parameter estimates in the presence of long-term nonstationarity in the noise spectrum. Signals used in this example will be in the Advanced LIGO band for  $\sim 100$  s. As shown in Figs. 5 and 6 the impact of nonstationary noise on parameter

estimation is likely to be more substantial as longer durations of data are needed for the analysis and spectral estimation. Quantifying how BayesLine improves parameter estimation for Advanced detector data is beyond the scope of this paper but our preliminary results are promising. We will embark on a large scale study using S6 data “recolored” to Advanced LIGO sensitivities and a broader range of inspiral signals to fully address the role BayesLine has to play in parameter estimation of future detections.

BayesLine has been developed outside of any existing LIGO/Virgo analysis pipelines and, for this work, is used as a stand-alone tool for PSD estimation. We have described the algorithm and demonstrated its ability to produce estimates to LIGO noise spectra which offer advantages over methods historically used for binary in-spiral parameter estimation analyses. BayesLine only addresses the stationary, Gaussian component of the noise and is otherwise insensitive to glitches which have caused problems for the analysis both in regards to false alarm rates and parameter estimation biases. A companion algorithm which does model transient non-Gaussian noise has been combined with BayesLine and is described in Ref. [33].

The BayesLine method will reach its full potential when it is integrated into other data analysis pipelines. We will next look to incorporate our spectral estimation method into the LALInference pipeline for parameter estimation follow up of binary in-spiral and merger signals.

The priors used in BayesLine are very wide and pose challenges for computing evidence integrals. A high priority improvement is to develop more informative priors which will improve efficiency of the RJMCMC and make evidence integrals tractable. One method currently under consideration is to construct kernel density estimates of the PSD using archived data. In this paradigm, BayesLine would run continuously on data as it is collected and continually update the kernel density estimate used for the prior in follow-up analysis of candidate signals.

BayesLine still implicitly assumes that the noise in the segment of data being analyzed is stationary which we have demonstrated not to be true for the low-mass compact binary signals. We envision developing a version of BayesLine that works across multiple small segments, with the parameters in each segment linked together by using the spectrum from neighboring segments as a prior.

Our model for the PSD is effective but is completely phenomenological. A physically motivated model for the broadband noise and spectral lines will provide more information about the detector. For the smooth broadband contribution to the PSD a physical model for the quantum noise can be found in Ref. [56] and phenomenological models for seismic and thermal noise have been constructed—all of which are parametrized by quantities relating to the detector. A more physically motivated model for the lines is also an intriguing possibility. For example, the 60 Hz line invites coherent modeling and subtraction

from the data, as opposed to whitening. Doing so will provide a cleaner residual for long-duration data segments and could also provide improved efficiency because harmonics may not need to be independently modeled. We conducted a preliminary study when deciding between regressing and whitening the spectral lines the results of which can be found in the Appendix.

### ACKNOWLEDGMENTS

We would like to acknowledge hospitality of the Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA) at Northwestern University where the BayesLine algorithm was, in part, conceived. Will Farr and Vicky Kalogera participated in the early discussions that resulted in the BayesLine algorithm and have provided suggestions throughout its development. We thank Michael Coughlin, Yiming Hu, and Matthew Pitkin for helpful comments on an earlier draft of this paper. N.J.C. appreciates the support of NSF Award PHY-1306702. T. B. L. acknowledges the support of NSF Award PHY-1307020.

### APPENDIX: SUBTRACTING THE POWER LINES

An alternative to the whitening approach taken by BayesLine when accounting for spectral lines is to coherently model and regress spectral lines from the data. Instead of modeling the *power* contained in the spectral lines—appearing in the denominator of Eq. (1)—we can include the time-dependent *phase* and *amplitude* evolution of the spectral lines in our calculation of the residual in the numerator, i.e. the residual becomes

$$\tilde{r}(f; \theta) = \tilde{d}(f) - \tilde{h}(f; \theta) - \tilde{l}(f; \lambda) \quad (\text{A1})$$

where  $\tilde{l}(f; \lambda)$  is a model for the spectral lines.

We will demonstrate this approach by focusing on the 60 Hz power line found in the LIGO data. Our example will use a 1024 s segment of data from the LIGO Livingston Observatory taken during S6. We choose the 60 Hz line because a Lorentzian is less physically motivated as the line profile—it is not a damped driven harmonic oscillator.

If the 60 Hz line had constant frequency it would put in a single spike in the Fourier series of the data and would be well approximated by a sinusoid. The line width spreads due to couplings between the electronics and the strain measurement and finite sampling effects. Temporal evolution will further broaden the spectral lines in *a priori* unpredictable ways. Sinusoids are a convenient basis to work in to model the noise at 60 Hz, but the data will demand another adaptive-dimension linear combination as is done with the Lorentzian-based model for the spectral lines discussed in this paper.

We experimented with modeling the power lines as sinusoids with amplitude  $A(t)$  and frequency  $f(t)$  which we




FIG. 9 (color online). The power spectrum of the data [red, solid], the coherent 60 Hz line model [green, dashed] and the residual [blue, dotted] after the line is subtracted.

compute in the time-domain. A Fast Fourier Transform (FFT) of the time-domain model is used to determine  $\tilde{l}(f; \lambda)$  which is then subtracted from the frequency-domain data as in Eq. (A1). The functions  $A(t)$  and  $f(t)$  are modeled by a linear combination of sines and cosines, with the number of basis functions used in each model determined by a RJMCMC.

Figure 9 shows the power spectrum of the data [red, solid], the coherent 60 Hz line model [green, dashed] and the residual [blue, dotted] after the line is subtracted. We found  $f(t)$  required  $\sim 60$  basis functions, while  $A(t)$  was easier to model, using only  $\sim 10$ . We find this to be an intriguing result and plan to further explore the benefits of coherently modeling spectral lines instead of whitening.




FIG. 10 (color online). The frequency [left axis; red, solid] and amplitude [right axis; blue short-dashed] of the data as a function of time as estimated by taking FFTs of each 16 s interval. The best fit model for  $f(t)$  and  $A(t)$  derived from the full 1024 s of data are shown in the green long-dashed curve and the magenta dotted curve, respectively.

To compare our time-domain model with the real data, we estimate the central frequency and peak amplitude of the 60 Hz line by subdividing the data into 16 s intervals, Fourier transforming each sub-segment of data, and finding the frequency with peak Fourier power. Figure 10 shows the peak frequency [left axis; red, solid] and peak amplitude (square root of peak Fourier power) [right axis; blue short-dashed] as a function of time for each 16 second interval. Our best fit model for  $f(t)$  and  $A(t)$  found from analyzing the full 1024 seconds of data using the sinusoid RJMCMC model is shown in the green long-dashed curve and the magenta dotted curve, respectively. We find the agreement between the time-evolving peak frequency and our model to be remarkable. These are two totally independent ways of estimating the frequency evolution of the 60 Hz line.

Note that unlike the Lorentzian model, which does power spectrum whitening, the power line model does coherent subtraction. The line-subtraction approach can be adapted to simultaneously remove harmonics of the 60 Hz line using a single model for  $f(t)$  and  $A(t)$  plus an overall scale factor for the amplitude of each harmonic. Early indications of this approach were promising but more work is necessary for this approach to reach maturity

The multi-harmonic line model could potentially reduce the overall dimension of the line model which would in turn simplify the MCMC implementation.

The sum-of-sinusoids approach worked well for the power lines but when it was tested on other narrow-band features in the data, especially lines from the mirror suspension system, the model required a challenging number of basis functions. As discussed in the main text of this paper, the suspension lines are noise driven and damped harmonic oscillators which are better modeled by the Lorentzians. Conveniently, the Lorentzian model also doing a fine job with the power and calibration lines. For our first implementation of BayesLine we therefore elected to use Lorentzians for all narrow-band features rather than have a complicated multicomponent model for the different types of lines.

We find the results of the sinusoid model to be of great interest and will continue developing the model for the power lines including higher harmonics. We will also investigate how incorporating environmental monitoring channels at the observatories may lead to more direct measurement, and therefore better modeling, of the time evolution of the electrical power supply and other narrow-band features.

- 
- [1] B. Abbott *et al.* (LIGO Scientific Collaboration), *Rep. Prog. Phys.* **72**, 076901 (2009).
  - [2] T. Accadia *et al.*, *JINST* **7**, P03012 (2012).
  - [3] J. Abadie *et al.* (LIGO Scientific Collaboration), *arXiv:1411.4547*.
  - [4] F. Acernese *et al.* (Virgo Collaboration), *Classical Quantum Gravity* **32**, 024001 (2015).
  - [5] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **80**, 102002 (2009).
  - [6] J. Abadie *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Phys. Rev. D* **82**, 102001 (2010).
  - [7] J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, C. Affeldt *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **85**, 122001 (2012).
  - [8] N. J. Cornish and J. D. Romano, *Phys. Rev. D* **87**, 122003 (2013).
  - [9] L. S. Finn, *Phys. Rev. D* **46**, 5236 (1992).
  - [10] C. Cutler, T. A. Apostolatos, L. Bildsten, L. S. Finn, E. E. Flanagan *et al.*, *Phys. Rev. Lett.* **70**, 2984 (1993).
  - [11] J. Aasi *et al.* (LIGO Collaboration, Virgo Collaboration), *Phys. Rev. D* **88**, 062001 (2013).
  - [12] J. Veitch, V. Raymond, B. Farr, W. Farr, P. Graff *et al.*, *Phys. Rev. D* **91**, 042003 (2015).
  - [13] P. Canitrot, *Phys. Rev. D* **63**, 082005 (2001).
  - [14] P. Ajith and S. Bose, *Phys. Rev. D* **79**, 084032 (2009).
  - [15] A. Buonanno, Y. Pan, J. G. Baker, J. Centrella, B. J. Kelly, S. T. McWilliams, and J. R. van Meter, *Phys. Rev. D* **76**, 104049 (2007).
  - [16] T. B. Littenberg, J. G. Baker, A. Buonanno, and B. J. Kelly, *Phys. Rev. D* **87**, 104003 (2013).
  - [17] M. Favata, *Phys. Rev. Lett.* **112**, 101101 (2014).
  - [18] N. Cornish, L. Sampson, N. Yunes, and F. Pretorius, *Phys. Rev. D* **84**, 062003 (2011).
  - [19] T. G. F. Li, W. Del Pozzo, S. Vitale, C. Van Den Broeck, M. Agathos, J. Veitch, K. Grover, T. Sidery, R. Sturani, and A. Vecchio, *Phys. Rev. D* **85**, 082003 (2012).
  - [20] L. Sampson, N. Cornish, and N. Yunes, *Phys. Rev. D* **89**, 064037 (2014).
  - [21] E. E. Flanagan and T. Hinderer, *Phys. Rev. D* **77**, 021502 (2008).
  - [22] J. S. Read, C. Markakis, M. Shibata, K. Uryu, J. D. E. Creighton, and J. L. Friedman, *Phys. Rev. D* **79**, 124033 (2009).
  - [23] B. D. Lackey, K. Kyutoku, M. Shibata, P. R. Brady, and J. L. Friedman, *Phys. Rev. D* **89**, 043009 (2014).
  - [24] L. Wade, J. D. E. Creighton, E. Ochsner, B. D. Lackey, B. F. Farr, T. B. Littenberg, and V. Raymond, *Phys. Rev. D* **89**, 103012 (2014).
  - [25] B. Allen, J. D. Creighton, E. E. Flanagan, and J. D. Romano, *Phys. Rev. D* **67**, 122002 (2003).
  - [26] C. Rover, R. Meyer, and N. Christensen, *Classical Quantum Gravity* **28**, 015010 (2011).

- [27] C. Rover, *Phys. Rev. D* **84**, 122004 (2011).
- [28] S. Vitale, G. Congedo, R. Dolesi, V. Ferroni, M. Hueller *et al.*, *Phys. Rev. D* **90**, 042003 (2014).
- [29] J. Veitch and A. Vecchio, *Phys. Rev. D* **78**, 022001 (2008).
- [30] J. Veitch and A. Vecchio, *Phys. Rev. D* **81**, 062003 (2010).
- [31] T. B. Littenberg and N. J. Cornish, *Phys. Rev. D* **82**, 103007 (2010).
- [32] T. B. Littenberg, M. Coughlin, B. Farr, and W. M. Farr, *Phys. Rev. D* **88**, 084044 (2013).
- [33] N. J. Cornish and T. B. Littenberg, [arXiv:1410.3835](https://arxiv.org/abs/1410.3835).
- [34] L. S. Finn, eConf C9808031, 07 (1998).
- [35] L. Finn, G. Gonzalez, J. Hough, M. Huq, S. Mohanty *et al.*, *AIP Conf. Proc.* **523**, 451 (2000).
- [36] S. Klimentenko, I. Yakushin, A. Mercer, and G. Mitselmakher, *Classical Quantum Gravity* **25**, 114029 (2008).
- [37] S. J. Waldman *et al.*, [arXiv:1103.2728](https://arxiv.org/abs/1103.2728).
- [38] G. L. Bretthorst, *Bayesian Spectrum Analysis and Parameter Estimation* (Springer, New York, USA, 1988).
- [39] P. Fritschel, R. Adikhari, and R. Weiss, Report No. IIGO-T050252, 2005.
- [40] Kawabe, *J. Phys. Conf. Ser.* **120**, 032003 (2008).
- [41] L. S. Finn and S. Mukherjee, *Phys. Rev. D* **63**, 062004 (2001).
- [42] T. Z. Summerscales *et al.*, Report No. IIGO-T060011-00-Z, 2006.
- [43] A. M. Sintes and B. F. Schutz, *Phys. Rev. D* **58**, 122003 (1998).
- [44] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. Teller, and E. Teller, *J. Chem. Phys.* **21**, 1087 (1953).
- [45] W. Hastings, *Biometrika* **57**, 97 (1970).
- [46] N. J. Cornish and J. Crowder, *Phys. Rev. D* **72**, 043005 (2005).
- [47] M. V. van der Sluys, C. Röver, A. Stroeer V. Raymond, I. Mandel, N. Christensen, V. Kalogera, R. Meyer, and A. Vecchio, *Astrophys. J.* **688**, L61 (2008).
- [48] B. Gough, *GNU Scientific Library Reference Manual - Third Edition* (Network Theory Ltd., 2009), 3rd ed., ISBN 0954612078, 9780954612078.
- [49] P. Green, *Highly Structured Stochastic Systems* (Oxford University Press, New York, 2003).
- [50] R. Umstätter, N. Christensen, M. Hendry, R. Meyer, V. Simha, J. Veitch, S. Vigeland, and G. Woan, *Phys. Rev. D* **72**, 022001 (2005).
- [51] N. J. Cornish and T. B. Littenberg, *Phys. Rev. D* **76**, 083006 (2007).
- [52] T. B. Littenberg and N. J. Cornish, *Phys. Rev. D* **80**, 063007 (2009).
- [53] A. Stroeer and J. Veitch, *Phys. Rev. D* **80**, 064032 (2009).
- [54] L. Sampson, N. Cornish, and N. Yunes, *Phys. Rev. D* **87**, 102001 (2013).
- [55] W. M. Farr, N. Sravan, A. Cantrell, L. Kreidberg, C. D. Bailyn, I. Mandel, and V. Kalogera, *Astrophys. J.* **741**, 103 (2011).
- [56] A. Buonanno and Y.-b. Chen, *Phys. Rev. D* **64**, 042006 (2001).