





Fast Parameter Estimation of Binary Mergers for Multimessenger Follow-up

Daniel Finstad  and Duncan A. Brown 

Department of Physics, Syracuse University, Syracuse, NY 13244, USA; dfinstad@syr.edu

Received 2020 October 7; revised 2020 November 12; accepted 2020 November 13; published 2020 December 10

Abstract

Significant human and observational resources have been dedicated to electromagnetic follow-up of gravitational-wave events detected by Advanced LIGO and Virgo. As the sensitivity of LIGO and Virgo improves, the rate of sources detected will increase. Margalit & Metzger (2019) have suggested that it may be necessary to prioritize observations of future events. Optimal prioritization requires a rapid measurement of a gravitational-wave event’s masses and spins, as these can determine the nature of any electromagnetic emission. We extend the relative binning method of Cornish (2013) and Zackay et al. (2018) to a coherent detector-network statistic. We show that the method can be seeded from the output of a matched-filter search and used in a Bayesian parameter measurement framework to produce marginalized posterior probability densities for the source’s parameters within 20 minutes of detection on 32 CPU cores. We demonstrate that this algorithm produces unbiased estimates of the parameters with the same accuracy as running parameter estimation using the standard gravitational-wave likelihood. We encourage the adoption of this method in future LIGO–Virgo observing runs to allow fast dissemination of the parameters of detected events so that the observing community can make best use of its resources.

Unified Astronomy Thesaurus concepts: [Gravitational waves \(678\)](#); [Gravitational wave astronomy \(675\)](#); [Astronomical methods \(1043\)](#); [Close binary stars \(254\)](#); [Compact binary stars \(283\)](#); [Neutron stars \(1108\)](#); [Astronomy data analysis \(1858\)](#)

1. Introduction

The observation of the binary neutron star merger GW170817 in gravitational and electromagnetic waves (Abbott et al. 2017b, 2017c) has demonstrated the importance of multimessenger astronomy in answering fundamental questions in physics, astronomy, and cosmology; see, e.g., Abbott et al. (2017a, 2017d) and Lattimer (2019). With the observation of GW190814, gravitational-wave astronomy has begun to explore the properties of compact objects that are more massive than previously observed neutron stars and less massive than previously observed black holes (Abbott et al. 2020). Advanced LIGO and Virgo perform a search for compact-object binary mergers with several low-latency analyses based on matched filtering (Messick et al. 2017; Nitz et al. 2018) and release alerts to the astronomical community to enable follow-up of detected events. As the sensitivity of the Advanced LIGO and Virgo detectors improves, the rate at which interesting events are detected will increase. Margalit & Metzger (2019) have suggested that it may become necessary to prioritize events for follow-up in future LIGO–Virgo observing runs. Optimal prioritization will require the knowledge of the source-frame component masses and spins of the binary, as these determine the type of electromagnetic counterpart that may be generated by the merger (Foucart et al. 2018; Capano et al. 2020).

In this Letter, we demonstrate that it is possible to perform full Bayesian parameter estimation on binary neutron star and neutron star–black hole signals within 20 minutes of the source’s detection by a matched-filter search (with an average time of 10.8 minutes) using 32 CPU cores (2.3 GHz Xeon® Gold 6140). Our analysis produces marginalized posterior probability densities for the source’s parameters (including source-frame masses, spins, sky location, and distance) that can be used to guide the prioritization of electromagnetic follow-up

in future LIGO–Virgo observing runs. We achieve this by extending the relative binning method originally introduced by Cornish (2013) (and independently developed by Zackay et al. 2018) to a fully coherent statistic, seeding the relative binning algorithm from the output of a matched-filter search, and using the *dynesty* nested-sampling package (Speagle 2020). We have made our code available in the PyCBC Inference framework (Biwer et al. 2019).

We validate our analysis on a population of simulated binary neutron star and neutron star–black hole signals in a LIGO–Virgo detector network. A matched-filter search is used to identify signals that have a false-alarm rate better than one per month. We then use our algorithm to produce marginalized posterior probability densities for each qualifying signal. For the parameters of interest, we perform a percentile–percentile test and demonstrate that our method produces unbiased parameter estimates. Comparing our sky localization to that of the *Bayestar* algorithm (Singer & Price 2016), we find that the 90% credible localization area improves by an average of 14 deg². We find that our analysis can recover the source-frame chirp mass to an accuracy of $\sim 5 \times 10^{-2} M_{\odot}$ for binary neutron star signals and $\sim 10^{-1} M_{\odot}$ for neutron star–black hole signals. We demonstrate that the measurement of mass ratio and spin is consistent with that of parameter estimation using the full likelihood, although these quantities are measured less accurately than the chirp mass as they enter the gravitational waveform at higher order and suffer from a partial degeneracy (Cutler & Flanagan 1994; Hannam et al. 2013). As an example use case, we demonstrate that our method recovers essentially the same posterior probabilities for the parameters of GW170817 as the full likelihood calculation. Our method obtains marginalized posteriors for GW170817 in 20 minutes, compared to over three hours using the standard likelihood calculation.

This Letter is organized as follows: In Section 2 we describe our simulated search. Section 3 describes our parameter estimation analysis and our implementation of relative binning for a detector network. Section 4 present our results including analysis run times and parameter estimation accuracy. Finally, we contrast our results to current methods in Section 5.

2. Simulated Search

We simulate a three-detector network representing the LIGO Hanford, LIGO Livingston (Abbott et al. 2016; Buikema et al. 2020), and Virgo (Acernese et al. 2015) detectors. We generate two populations of simulated signals: 600 binary neutron star and 570 neutron star–black hole binaries. Each population is injected into a realization of 33 hr of simulated detector data, which is created by coloring Gaussian noise to the design power spectral density of each detector (Abbott et al. 2018). The simulated binary neutron star signals have their chirp mass drawn uniformly from the interval $[0.5, 3]M_{\odot}$ and mass ratio $q = m_1/m_2$ drawn uniformly from the interval $[1, 3]$, with constraints on the component masses so that $1 < m_{1,2}/M_{\odot} < 3$. The neutron star’s spins are restricted to be aligned with the orbital angular momentum and have dimensionless magnitude drawn uniformly from the interval $[-0.05, 0.05]$. The simulated neutron star–black hole signals have their chirp mass drawn uniformly from the interval $[0.5, 7]M_{\odot}$ and mass ratio drawn uniformly from the interval $[1, 10]$, with constraints on the component masses so that $1 < m_1/M_{\odot} < 10$ and $1 < m_2/M_{\odot} < 3$. Both component spins are restricted to be aligned with the orbital angular momentum, with the black hole spin dimensionless magnitude drawn uniformly from the interval $[-0.998, 0.998]$ and the neutron star spin dimensionless magnitude drawn uniformly from the interval $[-0.05, 0.05]$. This population of sources is chosen to cover the region in which it is expected that there will be neutron star disruption and an electromagnetic counterpart (Capano et al. 2020). Each set of simulated signals is uniformly distributed in sky location and follow a uniform-in-volume distance distribution with $d_L \in [10, 300]$ Mpc for binary neutron star signals and $d_L \in [10, 500]$ Mpc for neutron star–black hole signals. This corresponds to a signal population with single-detector signal-to-noise ratios of 1 to $O(100)$. Binary neutron star signals are simulated using the TaylorF2 waveform approximant (Dhurandhar & Sathyaprakash 1994; Droz et al. 1999; Blanchet et al. 1995; Faye et al. 2012). The neutron star–black hole signals are simulated using the IMRPhenomD approximant (Husa et al. 2016; Khan et al. 2016). For both populations, we set the tidal deformability of the neutron stars Λ to zero, as this does not have a significant effect on the parameters we are investigating in this Letter (Damour et al. 2012).

To simulate the output of the LIGO–Virgo searches, we run each set of simulated signals through the PyCBC search pipeline (Usman et al. 2016) configured to operate in a similar way to the PyCBC Live low-latency search used in the recent Advanced LIGO–Virgo observing runs (Dal Canton et al. 2020). This search uses matched filtering (Allen et al. 2012) with a template bank of gravitational waveforms designed to give at least a 97% match, measured by noise-weighted overlap, to any potential signal in the relevant parameter space (Dal Canton & Harry 2017; Abbott et al. 2019a). The bank is designed to catch potentially electromagnetically bright signals, and contains 315,325 waveforms. Template waveforms have

component masses spanning $[1, 30]M_{\odot}$ and dimensionless spin magnitudes in the range $[-1, 1]$, with the spin restricted to the direction of the orbital angular momentum. Templates in the bank are generated using the TaylorF2 approximant for signals with total mass $M = m_1 + m_2 < 4M_{\odot}$ (Faye et al. 2012), and with a reduced-order model of the SEOBNRv4 approximant otherwise (Bohé et al. 2017). Candidate triggers are required to be matched by the same template in at least two detectors in the network and with consistent phase, amplitude, and time of arrival given the network orientation and relative sensitivities between detectors (Nitz et al. 2017). The search pipeline provides best-fit template parameters for every trigger and measures the trigger’s statistical significance. The significance of a trigger is determined by the time-slide method and the pipeline computes a false-alarm rate for each trigger. We select the triggers that have a false-alarm rate more significant than 1 per month as candidate events for parameter estimation follow-up. This threshold was selected to be the same as that used to release low-latency events as public alerts for electromagnetic follow-up in the third LIGO–Virgo observing run (Abbott et al. 2019b) and corresponds to a network signal-to-noise ratio of approximately 8.3. Of the total injections made, 306 binary neutron star and 253 neutron star–black hole injections satisfied this threshold.

3. Parameter Estimation

We use *PyCBC Inference* (Biver et al. 2019) with the *dynesty* nested sampler (Speagle 2020) to perform Bayesian parameter estimation on candidate events from the search pipeline. In general, under the assumption of Gaussian noise characterized by a power spectrum $S(f)$, the likelihood of obtaining detector data d given the presence of a gravitational waveform $h(\theta)$ is

$$\mathcal{L}(d|\theta) \propto \exp\left[-\frac{1}{2}\langle d - h(\theta) | d - h(\theta) \rangle\right], \quad (1)$$

where

$$\langle a|b \rangle = 4\Re \int_{f_{\min}}^{f_{\max}} \frac{\tilde{a}^*(f)\tilde{b}(f)}{S(f)} df \quad (2)$$

is the noise-weighted inner product (Chernoff & Finn 1993; Finn & Chernoff 1993). In evaluating this likelihood, we can obtain estimates of the gravitational-wave parameters θ through the posterior probability distribution

$$p(\theta|d) \propto \mathcal{L}(d|\theta)p(\theta), \quad (3)$$

where $p(\theta)$ is the assumed prior probability distribution of the parameters. To calculate the likelihood, we use the relative binning method of Cornish (2013) and Zackay et al. (2018), which uses a linear interpolation across frequency samples over which the accumulated phase difference $\delta\phi$ between a fiducial waveform and nearby waveforms is less than a tunable threshold. This effectively downsamples the number of frequency points used to compute the likelihood, thereby speeding up the parameter estimation.

The implementation of relative binning used by Zackay et al. (2018) did not incorporate a coherent network detection statistic. We extend their method to include the extrinsic parameters that are needed to measure the sky location of an event: R.A. α , decl. δ , geocentric time of coalescence t_c ,




Figure 1. Left: the wall-clock time in minutes that it takes to perform the parameter estimation using the coherent relative binning likelihood and nested sampling on 32 cores of an Intel® Xeon® Gold 6140 processor running at a clock speed of 2.3 GHz as a function of the network signal-to-noise ratio of the maximum likelihood template. The average run-time for a single signal is 10.8 minutes, with the maximum run-time being 20 minutes for all signals. Increasing the number of cores does not significantly decrease the wall-clock run-time. The run-time shows a slight increase as a function of the signal-to-noise ratio, as expected given that signals with a larger signal-to-noise ratio have a more narrowly peaked likelihood. Right: the fraction of injections recovered within a credible interval plotted as a function of credible interval. Fidelity to the 1:1 diagonal line is an indication of probability being uniformly distributed across a given parameter’s posterior distribution and is a measure of the accuracy of this analysis at the population level. We find that all of the parameters of interest are estimated in an unbiased way by our parameter estimation method.

inclination angle ι , and gravitational-wave polarization angle ψ . These parameters are incorporated into the likelihood by projecting each template waveform onto the individual detectors in the network. A general frequency domain waveform template h as seen by a detector can be written as

$$h(f) = F_+(\alpha, \delta, \psi)h_+(f) + F_\times(\alpha, \delta, \psi)h_\times(f), \quad (4)$$

where $h_{+,\times}$ are the plus and cross polarizations of the waveform, and $F_{+,\times}$ are the detector antenna responses to the two polarizations (Anderson et al. 2001). The amplitude of the individual waveform polarizations depend on the inclination angle ι (Thorne 1987)

$$h_+ \propto \frac{1}{2}(1 + \cos^2 \iota), \quad (5)$$

$$h_\times \propto \cos \iota. \quad (6)$$

We generate waveforms using both polarizations in order to capture this dependence. Similarly, we measure α , δ , t_c , and ψ dependence through the detector antenna responses as the orientation of the detector arms, and thus the sensitivity to the two polarizations, will change as the Earth moves. To account for coherent network timing delays, we calculate detector-specific arrival times for each template waveform using α , δ , and t_c , based on the geometry of the network with respect to the source at the time of the signal, along with the light travel time from the Earth center (Fairhurst 2009).

The relative-binned likelihood calculation requires a fiducial waveform known to be near the peak of the likelihood. The chirp mass of the template used to generate a candidate by a search pipeline is accurate to within a few $10^{-3} M_\odot$ for binary neutron star signals (Berry et al. 2016; Biscoveanu et al. 2019) and to approximately 1% for neutron star–black hole signals (Canton et al. 2020). Since the chirp mass is the leading order




Figure 2. Fraction of injections recovered as a function of the 90% confidence localization area. The localization results from our parameter estimation analysis are shown in blue, and those from the Bayestar algorithm are in orange. The dotted lines show the results for the entire set of signals, while solid lines show only signals that were above the detection threshold in all three detectors in our simulated search. We find our localization areas are consistently smaller than those from Bayestar, as indicated by the blue lines lying to the left of the orange lines, although the difference in areas is not large. The improvement in localization area between Bayestar and our analysis is 14 deg² on average, and is comparable between the set of triple-coincident signals and the set of all signals.

parameter governing phase evolution for a binary inspiral (Peters & Mathews 1963), the best-fit template will be near the peak of the likelihood. We therefore use the parameters that the search pipeline reports for a signal to generate the fiducial waveform that seeds the relative binning method. For the fiducial sky location, inclination, and polarization, we arbitrarily choose $\alpha_f = \pi$, $\delta_f = 0$, $l_f = 0$, and $\psi_f = \pi$, as we find

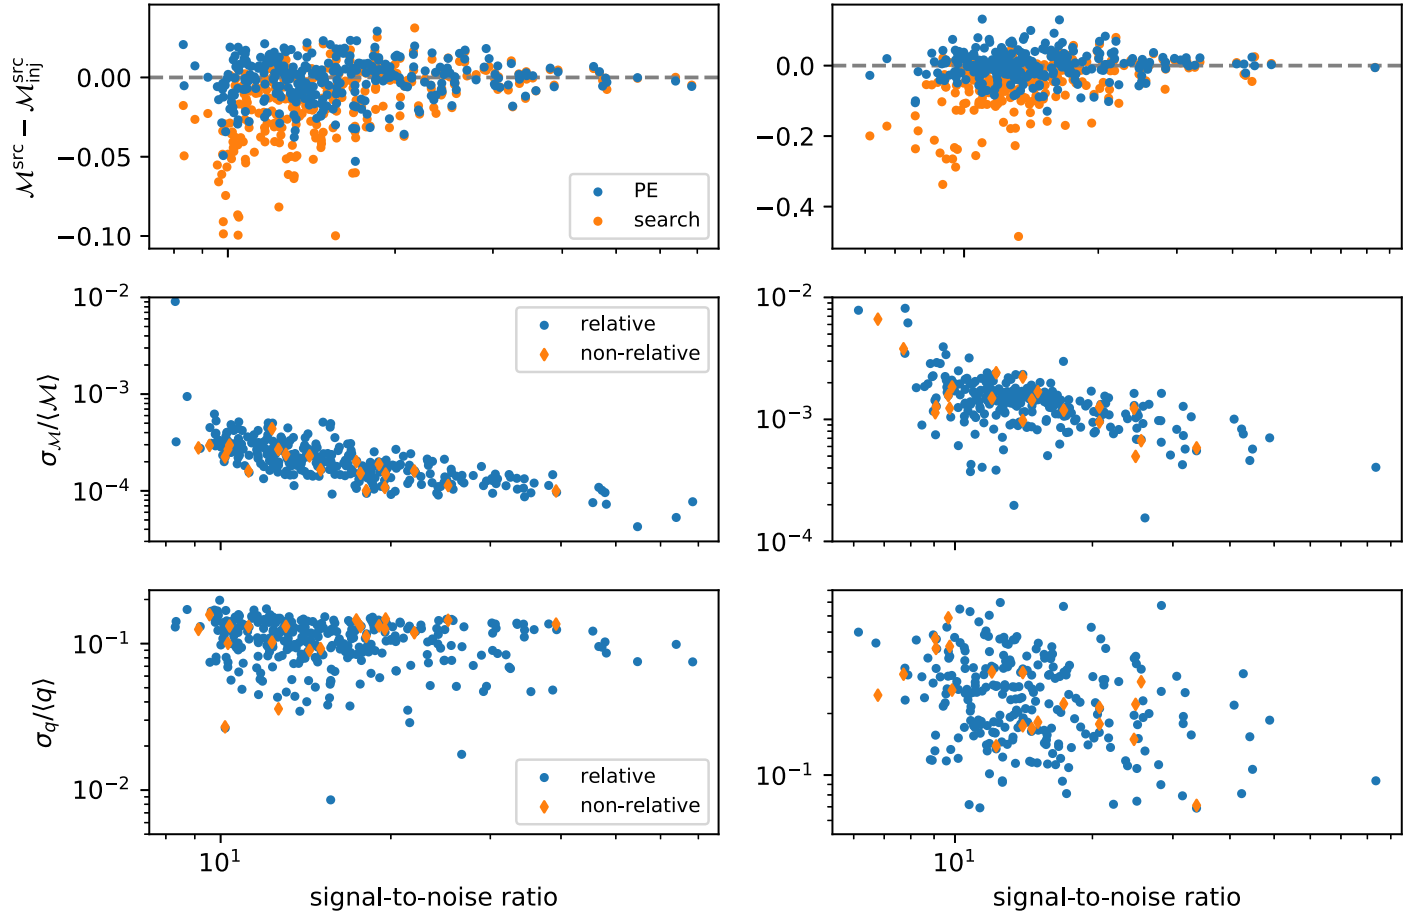


Figure 3. Chirp mass and mass ratio recovery metrics for the binary neutron star (left column) and neutron star–black hole (right column) signals in our analysis. Top row: difference between source-frame chirp mass estimates and the true injected value, as a function of signal-to-noise ratio. Blue circles denote differences from the median posterior values from parameter estimation, while orange circles show differences from best-fit template values by a factor of 2. Middle and bottom rows: fractional uncertainties on chirp mass and mass ratio, respectively, calculated as the ratio of standard deviation and mean of the posterior distributions. Uncertainties from our relative-binned analysis are shown as blue circles, and those from a standard nonrelative likelihood analysis on a subset of the population are shown as orange diamonds. Our relative-binned results are consistent with the nonrelative analysis, and also with the results in Farr et al. (2016).

that more accurate initial estimates are unnecessary to correctly recover the source parameters. The fiducial coalescence time is set to be the arithmetic mean of the coalescence time reported by the search pipeline for each detector.

Parameter estimation is performed over the detector-frame chirp mass \mathcal{M} , the mass ratio $q = m_1/m_2$, $m_1 \geq m_2$, the component-aligned spins $\chi_{1,2}$, the geocentric time of coalescence t_c , the inclination angle ι , the R.A. α , the decl. δ , the luminosity distance d_L , and the gravitational-wave polarization angle ψ . The likelihood calculation includes an analytic marginalization over the coalescence phase ϕ_c . We use the TaylorF2 approximant to generate the likelihood for binary neutron star waveforms and the IMRPhenomD approximant for the neutron star–black hole waveforms. For all simulated signals we use a low-frequency cutoff of 30 Hz and a sample rate of 2048 Hz.

The prior distributions used in the parameter estimation are the same as those of the corresponding population of simulated signals for each parameter, with the exception of the chirp mass that we restrict to be uniform in $\mathcal{M} \in [\mathcal{M}_s - 0.1, \mathcal{M}_s + 0.1] M_\odot$, where \mathcal{M}_s is the chirp mass of the template reported by the search. This constraint on the chirp mass prior enables quicker convergence of the parameter estimation, but in all

cases the restricted bounds are well outside the region of posterior support and so do not affect the accuracy of recovery.

For each simulated signal recovered with false-alarm rate more significant than 1 per month by the search pipeline, we run the relative-binned parameter estimation analysis to produce posterior distributions for the 10-dimensional set of waveform parameters $\theta = (\mathcal{M}, q, \chi_1, \chi_2, t_c, \iota, \alpha, \delta, d_L, \psi)$. For each signal, we measure the wall-clock time that it takes to perform the parameter estimation on 32 cores of an Intel® Xeon® Gold 6140 processor running at a clock speed of 2.3 GHz.

4. Results

The timing results for the two simulated populations as a function of the network signal-to-noise ratio of the maximum likelihood template are shown in the left panel of Figure 1. The average run time for a single signal is 10.8 minutes, with the maximum run-time being 20 minutes for all signals. The parallelization used by the nested-sampling algorithm is saturated at approximately 32 cores, so while a small decrease in wall-clock time may be gained by fine-tuning the number of cores, increasing the number beyond 32 does not significantly




Figure 4. Difference between parameter estimates and true injected values for some component parameters of interest, plotted against signal-to-noise ratio. The left column shows results for the component masses of binary neutron star signals, and the right column shows results for the component masses and black hole spin of neutron star–black hole signals. Differences are computed from median posterior values, and masses have been converted to the source frame using the distance posteriors. We find both component masses of binary neutron star signals are generally constrained to within $\sim 0.5 M_{\odot}$ of the true value for all signals, while the majority of primary and secondary masses of neutron star–black hole signals are within about $3 M_{\odot}$ and $1 M_{\odot}$, respectively. We find our black hole spin measurements are uninformative below a signal to noise of 20, but for louder signals the spin is within about 0.3 of the true value.

decrease the run time. Processors with a faster clock speed will generally decrease run time, however.

To determine whether our method of measuring the parameters is accurate for the population of injected signals, we perform a percentile–percentile (PP) test on each of the main parameters of interest: chirp mass \mathcal{M} , mass ratio q , effective spin $\chi_{\text{eff}} = (m_1\chi_1 + m_2\chi_2)/(m_1 + m_2)$, R.A. α , decl. δ , luminosity distance d_L , and inclination i . The PP test calculates the distribution of percentile ranks for all injected parameter values within their respective posteriors and constructs the fraction of injections recovered within a credible interval as a function of credible interval. Any deviation from uniformity in this distribution for a parameter is an indication of measurement bias. We measure any deviation with the Kolmogorov–Smirnov (KS) test (Massey 1951), which computes the distance between the empirical distribution that we find for the PP test and the expected distribution. The results of the PP tests are shown in the right panel of Figure 1. For every parameter of interest, we find that the PP test follows the ideal distribution well, with the KS test indicating that the percentile rank distributions cannot be meaningfully distinguished from uniform. Our results show that our analysis produces unbiased estimates for each of the parameters of interest.

To examine the accuracy of sky localization, we calculate the area on the sky containing 90% of the probability for the location of the source. We compare the area of this probability contour to the 90% credible interval of the sky-map produced in low latency by the *Bayestar* algorithm (Singer & Price 2016). Figure 2 shows the cumulative fraction of signals

recovered as a function of the 90% confidence localization area for our method and by *Bayestar*. For direct comparison to the results of Singer & Price (2016), we calculate the cumulative fraction using all recovered signals, and the subset of the recovered signals that is detected above threshold in all three detectors. We find that the area of the 90% credible region improves by an average of 14 deg^2 when using the relative binning parameter estimation compared to *Bayestar*.

To examine the accuracy of parameter recovery, we calculate the difference between the median of the posterior and the known injected value for each parameter. The accuracy of chirp mass recovery in the source frame is shown in the top panels of Figure 3 as a function of the network signal-to-noise ratio for each recovered signal. As expected, the accuracy of recovery increases as the signal to noise increases. For binary neutron star signals the difference between the median value of the chirp mass posterior and the injected value is less than $\sim 5 \times 10^{-2} M_{\odot}$ for all simulated signals. This accuracy improves by a factor of 2 for signal to noise greater than 20. Neutron star–black hole signals generally have larger uncertainties on their parameters and we find chirp mass residuals for these signals to be less than $10^{-1} M_{\odot}$ for signal to noise greater than 10 and a factor of 2 less than that for signal to noise greater than 20. For comparison, we show the accuracy of the source-frame chirp mass of the best-fit template from the search. The search measures the detector-frame parameters of the gravitational-wave signal, so we convert this to the source frame by computing the redshift at the median distance reported by *Bayestar* for the candidate event. The accuracy

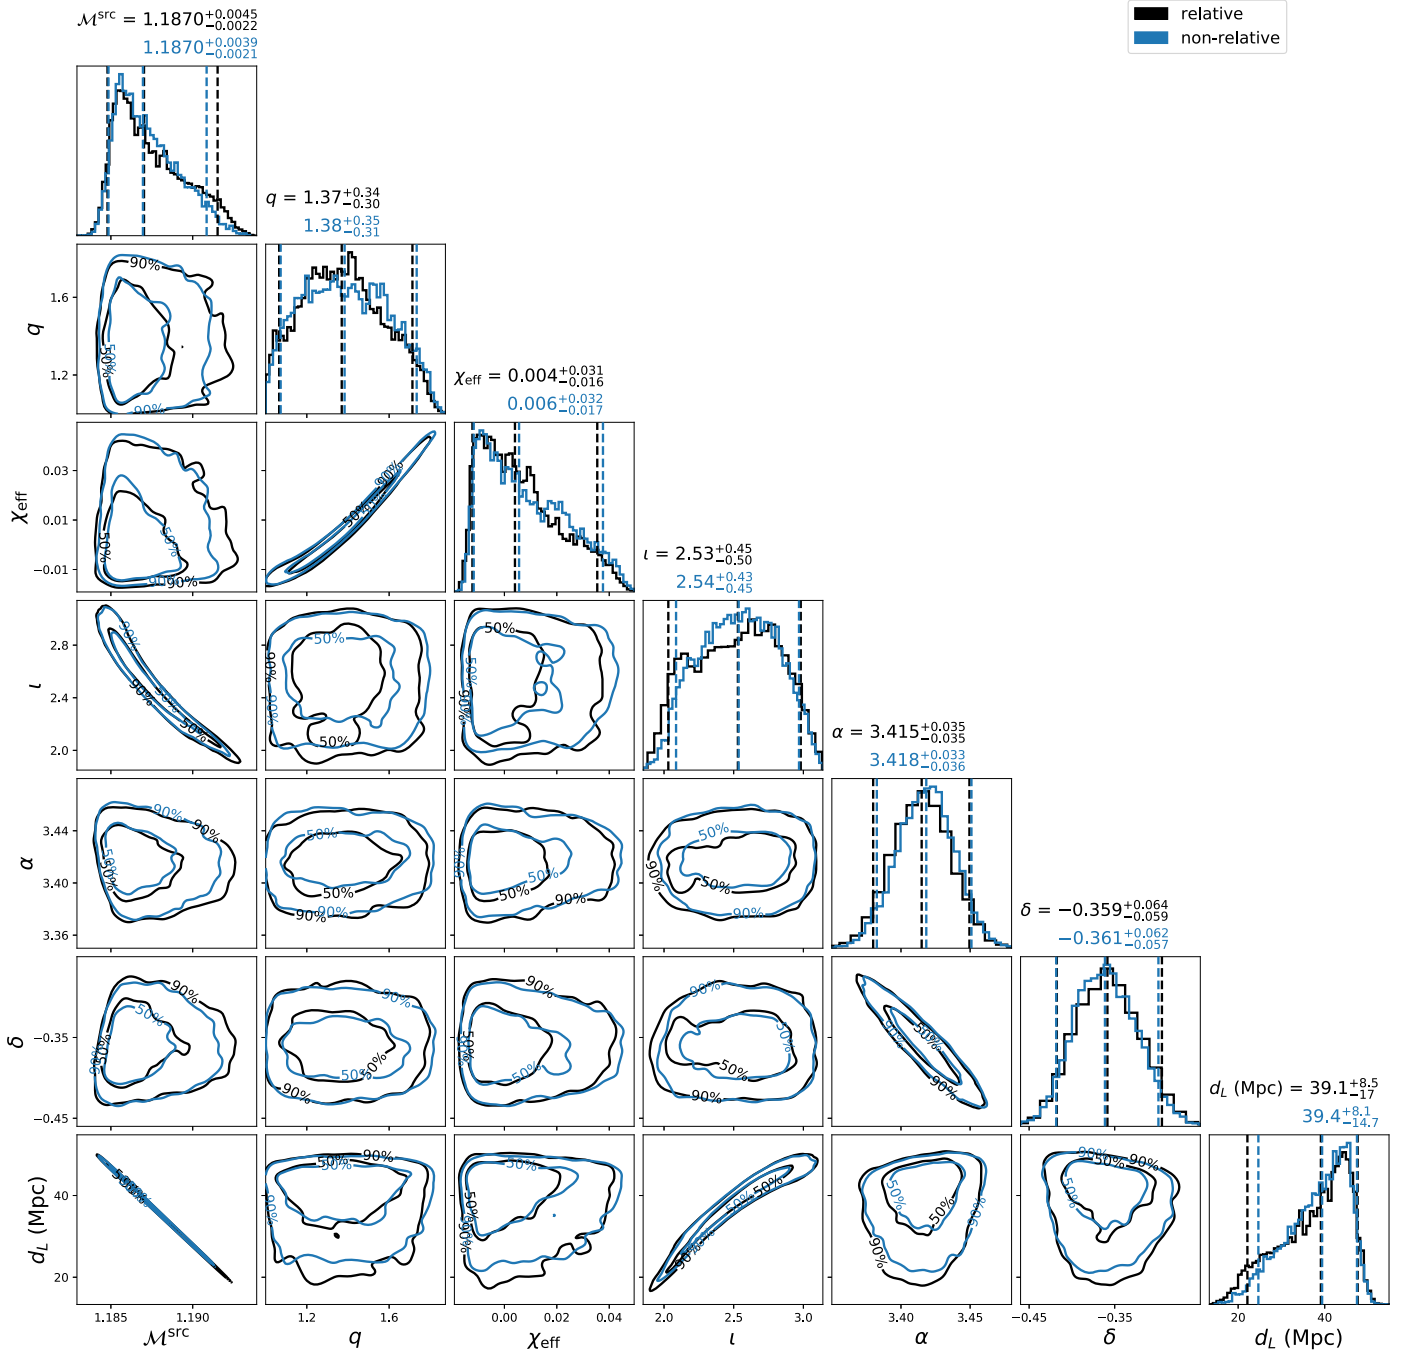


Figure 5. Posterior distributions from a relative-binned parameter estimation analysis of GW170817 (black contour) as compared to a run using the standard nonrelative likelihood (blue contour). Marginalized one-dimensional histograms for each parameter are shown along the diagonal, with vertical dashed lines at the median value and the bounds of the 90% credible interval. Off-diagonal plots show two-dimensional slices of the parameter space with contours delineating the 50% and 90% credible regions. The relative-binned analysis completed in 20 minutes versus roughly 3 hr in the nonrelative case, and all parameter distributions are consistent between the two analyses.

of the best-fit chirp mass from the search is an order of magnitude worse than estimated by Biscoveanu et al. (2019). However, the majority of the error comes from the calculation of the source-frame chirp mass. Comparing the detector-frame chirp mass of the simulated signal to the best-fit template, we find errors of $\sim 10^{-3} M_{\odot}$. While the accuracy of the best-fit template degrades for quieter signals, using this estimate as a seed for the relative-binned analysis does not affect the recovery of the source parameters.

The middle row of Figure 3 shows the fractional uncertainty in the chirp mass $\sigma_M / \langle M \rangle$, where σ_M and $\langle M \rangle$ are the standard deviation and mean of the posterior distribution. By this measure we find the accuracy of our method for the binary neutron star population is comparable to that of Farr et al. (2016). As an additional check we also run parameter estimation using the full likelihood for a subset of the population and find that the accuracy of the relative binning method is consistent with results using the full likelihood for both binary neutron star and neutron star–black hole signals.

These results show that our recovery of the chirp mass for all signals has more than sufficient accuracy to determine the expected type of electromagnetic counterpart and the possible fate of the merger remnant using the method of Margalit & Metzger (2019).

The gravitational-wave phase evolution is less sensitive to changes in the mass ratio and so the component masses of the binary are less well recovered than the chirp mass (Cutler & Flanagan 1994). A degeneracy exists between the mass ratio and component spins of the binary that makes measuring the component masses and spins challenging, especially for neutron star–black hole systems (Hannam et al. 2013). The bottom row of Figure 3 shows the accuracy of measuring the mass ratio $q = m_1/m_2$. Although the measurement of this parameter is less accurate than that of the chirp mass, our results are again comparable to those seen by Farr et al. (2016) and they are consistent with our comparison analysis using the full likelihood on a subset of the population. This demonstrates that the reduced accuracy is intrinsic to the measurability of the parameter and not a result of using the relative binning algorithm.

To further illustrate the utility of our method in recovering parameters of interest to the observing community, Figure 4 shows the source-frame component mass residuals for all signals as well as the black hole spin residuals for the neutron star–black hole signals, plotted as a function of the network signal-to-noise ratio. The binary neutron star component masses are shown in the left panels of the figure. The residuals on the primary mass are generally less than about $0.5 M_\odot$ with only a slight tendency to smaller values as signal to noise increases. The secondary mass residuals are somewhat smaller, less than about $0.3 M_\odot$, which can be attributed to the relatively narrow mass parameter space ($1\text{--}3 M_\odot$) and our convention requiring $m_2 < m_1$.

Neutron star–black hole signals have larger uncertainties on their intrinsic parameter estimates owing to the larger mass and spin parameter space and the known degeneracy between mass ratio and spin (Hannam et al. 2013). However, these quantities are important in determining whether a merger will produce an electromagnetic counterpart. The residuals on component masses and black hole spin for our neutron star–black hole signals are shown in the right panels of Figure 4. We find the primary and secondary mass residuals are mostly less than $3 M_\odot$ and $1 M_\odot$, respectively. Our estimates of the black hole spin are generally uninformative below a signal to noise of 20, but above this threshold we find the residuals are constrained to be less than ~ 0.3 .

As a final example of the effectiveness of our method, we apply it to GW170817 (Abbott et al. 2017b) without including any prior knowledge of host galaxy location or distance. For comparison, we also repeat the analysis using a standard nonrelative likelihood, and the posteriors from both runs are shown in Figure 5. For all measured parameters, we find the posterior distributions from the relative and nonrelative analyses are nearly identical, in agreement with Dai et al. (2018). However, the analysis using the relative binning likelihood seeded by a search took only 20 minutes to complete, as compared to over 3 hr for the standard likelihood computation. In the only confirmed observation of a multi-messenger gravitational-wave source to date, our analysis is able to provide the same localization region as the standard

likelihood as well as the same intrinsic parameter estimates in substantially less computational time.

5. Conclusion

In previous LIGO–Virgo observing runs, the information provided in low latency to astronomers consisted of the time of the signal, an estimate of its statistical significance (false-alarm rate), and a three-dimensional localization probability in sky location and distance. In the recent third observing run, two additional classifications were released that bin events into one of five broad categories (binary neutron star, binary black hole, neutron star–black hole, mass gap, or terrestrial noise) and estimate the probability that the event produced an electromagnetic counterpart (Canton et al. 2020; Kapadia et al. 2020). Both of these methods are based on the parameters of the best-fit matched-filter template recorded by the low-latency search. Biscoveanu et al. (2019) performed a template-bank simulation that estimated that the low-latency chirp mass point estimate for binary neutron stars is accurate to $\sim 10^{-3} M_\odot$; however, they note that there can be significant bias in mass ratio and effective spin from the best-fit template. Canton et al. (2020) demonstrated that the best-fit chirp mass from a search can be used to inform a classification scheme in which the classifications are correct in a large majority of cases.

Here, we have extended the relative binning algorithm (Cornish 2013; Zackay et al. 2018) for fast likelihood evaluation in gravitational-wave parameter estimation to a fully coherent detector network and demonstrated that it can be seeded by the output of a matched-filter search. We have applied our method to a set of 559 simulated signals (306 binary neutron star and 253 neutron star–black hole binaries) as well as to GW170817. We find that in all cases our method produces unbiased estimates for all measured parameters in less than 20 minutes. We have shown that our method is capable of producing full posterior distributions for all signal parameters, which do not suffer from the biases seen when attempting to measure the mass ratio and spin from the best-fit template. In the case of GW170817, the relative-binned analysis produces results nearly identical to those from a standard analysis using the full likelihood, emphasizing our method’s utility in producing fast parameter estimates that are of particular interest for electromagnetic follow-up.

For gravitational-wave events in LIGO’s third observing run, the average time between an initial trigger alert and the first Bayesian parameter estimation results being made available was about 10 hr (although only updated sky maps are released and not measurement of the source’s parameters). We have demonstrated our method could reduce this delay time considerably, which would allow for electromagnetic follow-up campaigns to be conducted more efficiently. We encourage the LIGO Scientific and Virgo collaborations to adopt these methods to provide the observing community with fast and accurate estimates of the parameters of detected signals so that these can be used to inform and prioritize electromagnetic follow-up strategies. Finally, we note that given the computational cost, very few large-scale injection studies of low-mass gravitational-wave signals have been done. Our implementation of the relative binning method into *PyCBC Inference* brings these sorts of studies within reach for even modestly equipped computing facilities.

We thank Brian Metzger, Ben Margalit, Edo Berger, and Tom Dent for helpful discussions. The authors are supported by National Science Foundation grant Nos. PHY-1707954 and PHY-2011655. D.A.B. thanks the Kavli Institute for Theoretical Physics for hospitality and partial support from National Science Foundation grant No. PHY-1748958. D.F. acknowledges support from a Syracuse University Research Excellence Doctoral Fellowship. Computational work was supported by Syracuse University and National Science Foundation award OAC-1541396. Supporting data for this Letter is available from <https://github.com/sugwg/rapid-relbin-pe>.

ORCID iDs

Daniel Finstad  <https://orcid.org/0000-0002-8133-3557>

Duncan A. Brown  <https://orcid.org/0000-0002-9180-5765>

References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016, *PhRvL*, **116**, 131103
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017a, *ApJL*, **848**, L13
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2018, *LRR*, **21**, 3
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019a, *PhRvX*, **9**, 031040
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019b, LIGO/Virgo Public Alerts User Guide, <https://emfollow.docs.ligo.org/userguide/index.html>
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017b, *PhRvL*, **119**, 161101
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017c, *ApJL*, **848**, L12
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017d, *Natur*, **551**, 85
- Abbott, R., Abbott, T. D., Abraham, S., et al. 2020, *ApJL*, **896**, L44
- Acernese, F., Agathos, M., Agatsuma, K., et al. 2015, *CQGra*, **32**, 024001
- Allen, B., Anderson, W. G., Brady, P. R., Brown, D. A., & Creighton, J. D. E. 2012, *PhRvD*, **85**, 122006
- Anderson, W. G., Brady, P. R., Creighton, J. D. E., & Flanagan, E. E. 2001, *PhRvD*, **63**, 042003
- Berry, C. P. L., Farr, B., Farr, W. M., et al. 2016, *Journal of Physics: Conference Series*, **716**, 012031
- Biscoveanu, S., Vitale, S., & Haster, C.-J. 2019, *ApJL*, **884**, L32
- Biwer, C., Capano, C. D., De, S., et al. 2019, *PASP*, **131**, 024503
- Blanchet, L., Damour, T., Iyer, B. R., Will, C. M., & Wiseman, A. G. 1995, *PhRvL*, **74**, 3515
- Bohé, A., Shao, L., Taracchini, A., et al. 2017, *PhRvD*, **95**, 044028
- Buikema, A., Cahillane, C., Mansell, G. L., et al. 2020, *PhRvD*, **102**, 062003
- Canton, T. D., Nitz, A. H., Gadre, B., et al. 2020, Realtime Search for Compact Binary Mergers in Advanced LIGO and Virgo's Third Observing Run Using PyCBC Live, arXiv:2008.07494
- Capano, C. D., Tews, I., Brown, S. M., et al. 2020, *NatAs*, **4**, 625
- Chernoff, D. F., & Finn, L. S. 1993, *ApJL*, **411**, L5
- Cornish, N. J. 2013, Fast Fisher Matrices and Lazy Likelihoods, arXiv:1007.4820
- Cutler, C., & Flanagan, E. E. 1994, *PhRvD*, **49**, 2658
- Dai, L., Venumadhav, T., & Zackay, B. 2018, arXiv:1806.08793
- Dal Canton, T., & Harry, I. W. 2017, arXiv:1705.01845
- Dal Canton, T., Nitz, A. H., Gadre, B., et al. 2020, arXiv:2008.07494
- Damour, T., Nagar, A., & Villain, L. 2012, *PhRvD*, **85**, 123007
- Dhurandhar, S., & Sathyaprakash, B. 1994, *PhRvD*, **49**, 1707
- Droz, S., Knapp, D. J., Poisson, E., & Owen, B. J. 1999, *PhRvD*, **59**, 124016
- Fairhurst, S. 2009, *NJPh*, **11**, 123006
- Farr, B., Berry, C. P. L., Farr, W. M., et al. 2016, *ApJ*, **825**, 116
- Faye, G., Marsat, S., Blanchet, L., & Iyer, B. R. 2012, *CQGra*, **29**, 175004
- Finn, L. S., & Chernoff, D. F. 1993, *PhRvD*, **47**, 2198
- Foucart, F., Hinderer, T., & Nisanke, S. 2018, *PhRvD*, **98**, 081501
- Hannam, M., Brown, D. A., Fairhurst, S., Fryer, C. L., & Harry, I. W. 2013, *ApJ*, **766**, L14
- Husa, S., Khan, S., Hannam, M., et al. 2016, *PhRvD*, **93**, 044006
- Kapadia, S. J., Caudill, S., Creighton, J. D. E., et al. 2020, *CQGra*, **37**, 045007
- Khan, S., Husa, S., Hannam, M., et al. 2016, *PhRvD*, **93**, 044007
- Lattimer, J. 2019, *Annals Phys.*, **411**, 167963
- Margalit, B., & Metzger, B. D. 2019, *ApJL*, **880**, L15
- Massey, F. J. 1951, *J. Am. Stat. Assoc.*, **46**, 68
- Messick, C., Blackburn, K., Brady, P., et al. 2017, *PhRvD*, **95**, 042001
- Nitz, A. H., Dal Canton, T., Davis, D., & Reyes, S. 2018, *PhRvD*, **98**, 024050
- Nitz, A. H., Dent, T., Dal Canton, T., Fairhurst, S., & Brown, D. A. 2017, *ApJ*, **849**, 118
- Peters, P., & Mathews, J. 1963, *PhRv*, **131**, 435
- Singer, L. P., & Price, L. R. 2016, *PhRvD*, **93**, 024013
- Speagle, J. S. 2020, *MNRAS*, **493**, 3132
- Thorne, K. S. 1987, in *Three Hundred Years of Gravitation*, ed. S. W. Hawking & W. Israel (Cambridge: Cambridge Univ. Press), 330
- Usman, S. A., Nitz, A. H., Harry, I. W., et al. 2016, *CQGra*, **33**, 215004
- Zackay, B., Dai, L., & Venumadhav, T. 2018, arXiv:1806.08792