

Removing degeneracy and multimodality in gravitational wave source parameters

Javier Roulet^{1,*}, Seth Olsen², Jonathan Mushkin³, Tousif Islam^{4,5,6,1},
Tejaswi Venumadhav^{7,8}, Barak Zackay³, and Matias Zaldarriaga⁹

¹*Kavli Institute for Theoretical Physics, University of California at Santa Barbara, Santa Barbara, California 93106, USA*

²*Department of Physics, Princeton University, Princeton, New Jersey 08540, USA*

³*Department of Particle Physics and Astrophysics, Weizmann Institute of Science, Rehovot 76100, Israel*

⁴*Department of Physics, University of Massachusetts, Dartmouth, Massachusetts 02747, USA*

⁵*Department of Mathematics, University of Massachusetts, Dartmouth, Massachusetts 02747, USA*

⁶*Center for Scientific Computing and Visualization Research, University of Massachusetts, Dartmouth, Massachusetts 02747, USA*

⁷*Department of Physics, University of California at Santa Barbara, Santa Barbara, California 93106, USA*

⁸*International Centre for Theoretical Sciences, Tata Institute of Fundamental Research, Bangalore 560089, India*

⁹*School of Natural Sciences, Institute for Advanced Study, 1 Einstein Drive, Princeton, New Jersey 08540, USA*

 (Received 5 August 2022; accepted 14 November 2022; published 15 December 2022)

Quasicircular binary black hole mergers are described by 15 parameters, of which gravitational wave observations can typically constrain only ~ 10 independent combinations to varying degree. In this work, we devise coordinates that remove correlations, and disentangle well- and poorly-measured quantities. Additionally, we identify approximate discrete symmetries in the posterior as the primary cause of multimodality, and design a method to tackle this type of multimodality. The resulting posteriors have little structure and can be sampled efficiently and robustly. We provide a PYTHON package for parameter estimation, `cogwheel`, that implements these methods together with other algorithms for accelerating the inference process. One of the coordinates we introduce is a spin azimuth that is measured remarkably well in several events. We suggest this might be a sensitive indicator of orbital precession, and we anticipate that it will shed light on the occurrence of spin-orbit misalignment in nature.

DOI: [10.1103/PhysRevD.106.123015](https://doi.org/10.1103/PhysRevD.106.123015)

I. INTRODUCTION

Gravitational wave astronomy is advancing at a rapid pace, with over 100 signals from compact binary mergers detected since the advanced LIGO [1] and advanced Virgo [2] detectors became operational [3–13]. The rate of detections will keep accelerating as planned and ongoing hardware upgrades take place. The astrophysical interpretation of these detections requires measuring their source parameters, such as masses, spins, position and orientation.

At the same time, following recent developments in waveform modeling, state-of-the-art models of compact binary coalescences now incorporate higher-order harmonics and generically oriented spins [14–17]. These improvements have increased the diversity of waveform morphologies and the parameter space dimensionality, rendering computational cost a major hurdle for parameter estimation studies.

In this work we develop a new parametrization of the source properties, tailor-made for compact binary mergers, that removes most correlations and multimodalities encountered in their posterior distributions. This coordinate system emphasizes quantities that are best constrained by the data, as opposed to the physically motivated parameters used by external libraries to model waveforms. We derive the set of analytic, invertible, nonlinear transformations between the standard parameter space and this new coordinate system. Expressed in these coordinates, the posterior distributions exhibit much less structure, which makes it easier to sample them efficiently and robustly.

In general, under a coordinate change $x \rightarrow y$ the prior and posterior probability densities transform as

$$P(y) = P(x)|J|, \quad (1)$$

*jroulet@ucsb.edu

where

$$|J| = |\partial x / \partial y| \quad (2)$$

is the absolute value of the Jacobian determinant. This can make coordinate changes cumbersome if their Jacobian determinant is not easily computable. We define our coordinate transformation as a series of simple transformations, affecting few variables at a time, that have tractable Jacobian determinants. We will use the following properties of the Jacobian. First, they multiply under composition of transformations $x \rightarrow y \rightarrow z$,

$$\left| \frac{\partial x}{\partial z} \right| = \left| \frac{\partial x}{\partial y} \right| \left| \frac{\partial y}{\partial z} \right|. \quad (3)$$

Second, transformations of the form

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \rightarrow \begin{pmatrix} a(x_2)x_1 + b(x_2) \\ x_2 \end{pmatrix}, \quad (4)$$

where x_1, x_2 may be multidimensional, have

$$|J| = |a(x_2)|, \quad (5)$$

which frees us to use arbitrarily sophisticated $b(x_2)$ functions. In particular, if $a(x_2) = \pm 1$ the probability density is unchanged by the coordinate transformation in Eq. (4).

The remainder of the article is structured as follows. In Sec. II we review the likelihood function of gravitational wave source parameters. In Sec. III we introduce coordinates that remove common correlations. In Sec. IV we identify approximate symmetries responsible for discrete degeneracies, and present a novel method to remove this type of multimodality. In Sec. V we present our main results; Sec. VA assesses the performance improvement brought by our methods in terms of the accuracy of the recovered posterior and the computational cost, and Sec. VB shows that a particular spin azimuth can be measured surprisingly well. We conclude in Sec. VI with a summary and outlook. Appendix A discusses in further detail two of the approximate symmetries identified, and Appendix B acts as a reference sheet in which we summarize the coordinate system proposed.

II. LIKELIHOOD

In this section we briefly review the likelihood function $P(d|\theta)$, through which the data d constrain source parameters θ . We cast its approximate dependence on extrinsic parameters in terms of the signal's amplitude, phase and time of arrival at each detector, and we provide analytical expressions for these.

Under the approximation that the noise is stationary and Gaussian,

$$\log P(d|\theta) = -\frac{1}{2} \sum_{k \in \text{det}} 4 \int_0^\infty df \frac{|\tilde{d}_k(f) - \tilde{h}_k(f; \theta)|^2}{S_k(f)}, \quad (6)$$

where \tilde{d}_k , \tilde{h}_k , and S_k are the frequency-domain data, strain model, and one-sided noise power spectrum in the k th detector, respectively [18].

For quasicircular binaries, the detector strains \tilde{h}_k depend on 15 parameters: eight intrinsic (two masses m_1, m_2 and six components of the spin vectors χ_1, χ_2) and seven extrinsic (luminosity distance d_L , inclination ι , right ascension α , declination δ , polarization ψ , orbital phase ϕ_{ref} , and coalescence time t_c). We can understand the approximate dependence on these parameters by considering the waveform under the quadrupole, stationary-phase and post-Newtonian approximations,

$$\tilde{h}_k(f; \theta) \approx A(f) \frac{\mathcal{M}^{5/6}}{d_L} R_k(\iota, \hat{\mathbf{n}}, \psi) e^{i[2\phi_{\text{ref}} - 2\pi f t_k(\hat{\mathbf{n}}) + \Psi(f; \theta_{\text{int}})]}, \quad (7)$$

where the chirp mass

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}} \quad (8)$$

controls the waveform amplitude and phase to leading post-Newtonian order. The inclination, sky location and polarization enter through the detector response [19]

$$R_k = \frac{1 + \cos^2 \iota}{2} F_k^+(\hat{\mathbf{n}}, \psi) - i \cos \iota F_k^\times(\hat{\mathbf{n}}, \psi), \quad (9)$$

where $\hat{\mathbf{n}}$ is the line of sight and F_k^+, F_k^\times are the antenna factors of the k th detector [20]. The sky location additionally introduces an arrival time delay in each detector due to the finite speed c of gravitational waves, so there is a separate arrival time t_k at each detector k .

Altogether, to this level of approximation the extrinsic parameters only affect the amplitude, phase and time in each detector,

$$\tilde{h}_k(f; \theta) \approx a_k e^{i\varphi_k} e^{-i2\pi(f - \bar{f}_k)t_k} A(f) e^{i\Psi(f; \theta_{\text{int}})}, \quad (10)$$

with

$$a_k = \frac{\mathcal{M}^{5/6}}{d_L} |R_k(\iota, \hat{\mathbf{n}}, \psi)|, \quad (11)$$

$$\varphi_k = \arg R_k(\iota, \hat{\mathbf{n}}, \psi) + 2\phi_{\text{ref}} - 2\pi \bar{f}_k t_k, \quad (12)$$

$$t_k = t_c - \frac{\hat{\mathbf{n}} \cdot \mathbf{r}_k}{c}, \quad (13)$$

where \mathbf{r}_k is the detector location relative to the point of reference time, e.g., the center of Earth, and \bar{f}_k is a frequency scale defined below. In Eq. (12) we have

included a term dependent on t_k to make φ_k orthogonal to t_k : to linear order, it is possible to absorb a small shift in the time of arrival t_k (in the sense of leaving the observed waveform almost unchanged) by simultaneously changing the phase (via ϕ_{ref} , say) by $2\pi\bar{f}_k t_k$ [21,22], where \bar{f}_k is the first frequency moment, i.e.,

$$\bar{f}_k^n := \frac{\int_0^\infty df A^2(f) f^n / S_k(f)}{\int_0^\infty df A^2(f) / S_k(f)} \quad (14)$$

for $n = 1$. This way, we expect the measurements of a_k , φ_k , t_k in a detector to be uncorrelated. We can also estimate the uncertainties in each of these parameters, which scale inversely with the signal-to-noise ratio ρ_k ,

$$\Delta a_k \sim \frac{a_k}{\rho_k}, \quad (15)$$

$$\Delta \varphi_k \sim \frac{1}{\rho_k}, \quad (16)$$

$$\Delta t_k \sim \frac{1}{2\pi\sigma_f \rho_k}, \quad (17)$$

with $\sigma_f^2 = \bar{f}^2 - \bar{f}^2$ [19,21]. By convention, for each event we will sort the detectors by their signal-to-noise ratio ($\rho_{k_0} > \rho_{k_1} > \dots$), so the best measured quantities correspond to the ‘‘reference detector’’ k_0 .

III. MITIGATING DEGENERACY

Oriented by the previous discussion, in this section we introduce a system of coordinates in which some variables separately control the observables a_{k_0} , φ_{k_0} , t_{k_0} , $t_{k_1} - t_{k_0}$, which will typically be well-measured, and others explore degeneracies and will be poorly constrained. Some of these coordinates are already known in the literature but we still review them here for completeness. As we will see, the posterior distribution expressed in these coordinates is largely uncorrelated.

A. Amplitude at reference detector

We replace the luminosity distance by the so-called chirp distance [23], which controls the amplitude at the reference detector Eq. (11),

$$\begin{aligned} \hat{d}(d_L, \mathcal{M}, \iota, \alpha, \delta, \psi) &:= \frac{1}{a_{k_0}} \\ &= \frac{d_L}{\mathcal{M}^{5/6} |R_{k_0}|}. \end{aligned} \quad (18)$$

This is proportional to the luminosity distance, with a scale factor that depends on the response of the detector to the given ι , α , δ , ψ , and the intrinsic amplitude of the source

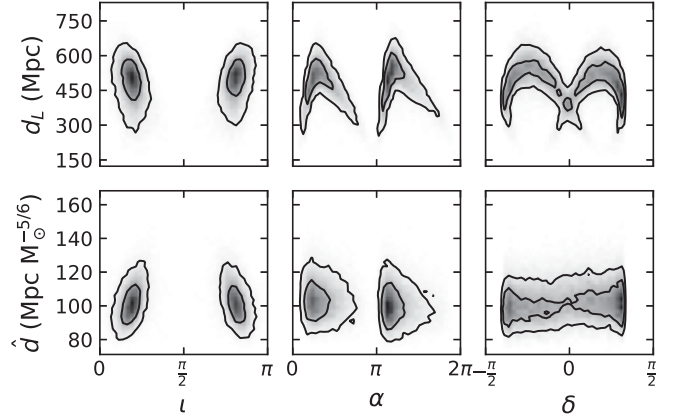


FIG. 1. The luminosity distance d_L is typically correlated with the inclination and sky location (top row), while \hat{d} [Eq. (18)] is less correlated and better measured (bottom row). Examples in this section correspond to GW151226 [24], but the highlighted qualitative features are generic.

$\propto \mathcal{M}^{5/6}$. Thus, \hat{d} avoids the correlations with these variables that the luminosity distance suffers, as shown in Fig. 1. A collateral benefit is that the observable values of \hat{d} are similar for all mass ranges (in contrast, heavier events can be observed at larger luminosity distances), so the distance range to explore does not need to be tuned for each event. From Eqs. (15) and (18) we expect that \hat{d} will typically be measured to a precision $\Delta \hat{d} \sim \hat{d} / \rho_{k_0}$.

In the notation of Eq. (4), the transformation $\hat{d} \rightarrow d_L$ has $x_1 = \hat{d}$, $x_2 = (\mathcal{M}, \iota, \alpha, \delta, \psi)$, $a(x_2) = 1 / \mathcal{M}^{5/6} |R_{k_0}|$ and $b = 0$. The Jacobian determinant is

$$|J| = \left| \frac{\partial \hat{d}}{\partial d_L} \right| = \frac{1}{\mathcal{M}^{5/6} |R_{k_0}|} = \frac{\hat{d}}{d_L}. \quad (19)$$

A more elaborate solution is to marginalize the posterior semianalytically over distance, altogether removing the necessity to sample from it [25,26]. We include this functionality as well in the software package we are releasing along with this paper. In practice, this makes the sampling process more robust against a particular failure mode, in which the sampler occasionally explores the distant Universe (favored by the prior) and misses a nearby solution with high likelihood.

B. Time of arrival at reference detector

We use the time of arrival at the reference detector t_{k_0} as our arrival time parameter [26], as opposed to, for example, the geocentric time of arrival.

From Eq. (17), the uncertainty in arrival time at the detector is typically $\lesssim 1$ ms [21]. On the other hand, unless the sky location is particularly well-constrained, the uncertainty in time of arrival at geocenter is on the order of the gravitational wave Earth-crossing time (roughly 40 ms) due




FIG. 2. The arrival time at geocenter (top row) is worse measured than at the reference detector (bottom row), because the time delay introduces a large correlation with sky location.

to the tight nonlinear correlation with sky location shown in Fig. 2 for the case of GW151226.

From Eqs. (4), (5), and (13), the transformation $t_{k_0} \rightarrow t_c$ has unit Jacobian determinant.

C. Time-of-arrival difference

When a signal is observed in multiple detectors, the arrival time differences provide the dominant constraint on its sky location. For each pair of detectors, the arrival time difference determines the angle $\theta_{kk'}$ between the source location and an axis through both detectors according to [see Eq. (13)]

$$\begin{aligned} t_k - t_{k'} &= \hat{\mathbf{n}} \cdot \frac{\mathbf{r}_{k'} - \mathbf{r}_k}{c} \\ &= \tau_{kk'} \cos \theta_{kk'}, \end{aligned} \quad (20)$$

where $\tau_{kk'}$ is the gravitational wave travel time between detectors k and k' .

Thus, a natural way of parametrizing the sky location measurement is with a polar coordinate system $(\theta_{\text{net}}, \phi_{\text{net}})$ whose z -axis contains the two detectors with the largest signal-to-noise ratios in the network [26,27], shown in Fig. 3. This coordinate system rotates with Earth and is related to the (fixed) α, δ by a 3D rotation that depends on the pair of detectors and the sidereal time at which the signal arrives. The isotropic prior is uniform in $\cos \theta_{\text{net}}, \phi_{\text{net}}$.

Figure 4 shows that, unlike α and δ , θ_{net} is typically well-measured and largely uncorrelated with ϕ_{net} . In fact, from Eqs. (17) and (20) its expected uncertainty is

$$\begin{aligned} \Delta \cos \theta_{\text{net}} &\sim \frac{1}{2\pi\sigma_f\tau_{k_0k_1}} \sqrt{\rho_{k_0}^{-2} + \rho_{k_1}^{-2}} \\ &= 0.016 \cdot \frac{100 \text{ Hz}}{\sigma_f} \cdot \frac{\tau_{\text{HL}}}{\tau_{k_0k_1}} \cdot \frac{10}{\rho_{k_1}} \sqrt{1 + \frac{\rho_{k_1}^2}{\rho_{k_0}^2}} \end{aligned} \quad (21)$$




FIG. 3. Coordinate system based on the dominant pair of detectors in the network, in this example Hanford-Livingston. The z -axis contains the two detectors and the y -axis points perpendicularly upwards. The zenith angle θ_{net} of the line-of-sight $\hat{\mathbf{n}}$ determines the time-of-arrival difference. (Detector sizes are exaggerated but otherwise the figure is to scale.)




FIG. 4. Sky location of events in the first two observing runs, expressed in terms of right ascension and declination (top) or network-based angles $\theta_{\text{net}}, \phi_{\text{net}}$ defined in Fig. 3 (bottom). θ_{net} is typically well constrained and largely uncorrelated with ϕ_{net} . The signals are clearly clustered near $\hat{\mathbf{n}} = \pm \hat{\mathbf{y}}_{\text{net}}$, at the two high-sensitivity lobes of the network antenna pattern. Accordingly, when two or less detectors dominate the information content, the azimuthal angle ϕ_{net} can be significantly informed by the antenna-pattern dependent prior.

for events observed in two detectors. Each pair of detectors constrains the source location to a narrow ring in the celestial sphere. This degeneracy is broken when the signal is prominent in three or more detectors. The distribution of ϕ_{net} is often multimodal, and for this reason we will ultimately use a modified azimuthal coordinate $\hat{\phi}_{\text{net}}$ instead (see Sec. IV A).

For single-detector events, a coordinate system aligned with the arms of the detector is more convenient.

D. Arrival phase

We now seek a parametrization of the form in Eq. (4) that expresses the well-measured arrival phase φ_{k_0} in terms of our new coordinates. As is clear from Eq. (12), such a reparametrization must involve ι , $\hat{\mathbf{n}}$, ψ , ϕ_{ref} , and t_{k_0} . To achieve this, we replace the coalescence phase ϕ_{ref} by another 2π -periodic coordinate

$$\hat{\phi}_{\text{ref}}(\phi_{\text{ref}}, \iota, \hat{\mathbf{n}}, \psi, t_{k_0}) := \phi_{\text{ref}} + \frac{\arg R_{k_0}(\iota, \hat{\mathbf{n}}, \psi) - 2\pi \bar{f}_{k_0}^{\text{ML}} t_{k_0} - \varphi_{k_0}^{\text{ML}}}{2} \quad (22)$$

$$\equiv \frac{\varphi_{k_0} - \varphi_{k_0}^{\text{ML}}}{2}, \quad (23)$$

where we choose the constants $\bar{f}_{k_0}^{\text{ML}}$ and $\varphi_{k_0}^{\text{ML}}$ as the first frequency moment and arrival phase at the dominant detector for the (approximate) maximum likelihood source configuration. $\hat{\phi}_{\text{ref}}$ can be interpreted as the deviation of the coalescence phase ϕ_{ref} from the value that would make the arrival phase at the reference detector equal to $\varphi_{k_0}^{\text{ML}}$. Therefore, its posterior distribution should have a peak near 0. The factor of $1/2$ in Eq. (23) reflects the fact that the phase of the dominant quadrupolar radiation advances by twice the angle under an azimuthal rotation of the source. We thus expect the $\hat{\phi}_{\text{ref}}$ posterior to have a second mode near π . By virtue of Eqs. (16) and (23), these modes are largely uncorrelated with all other parameters and have widths $\sim 1/(2\rho_{k_0})$. Radiation harmonics with odd values of m break the symmetry between these two modes. Figure 5 shows that $\hat{\phi}_{\text{ref}}$ is indeed much better measured and less correlated than ϕ_{ref} . The transformation $\phi_{\text{ref}} \rightarrow \hat{\phi}_{\text{ref}}$ has unit Jacobian.

Despite removing correlations, the definition of $\hat{\phi}_{\text{ref}}$ in Eq. (22) has two undesirable properties for sampling: its posterior is bimodal and, for waveforms with higher modes, discontinuous at the branch cut of $\arg R_{k_0}$. We will solve both problems in Sec. IV B.

E. Reference frequency

Source parameters that are not conserved throughout the inspiral and merger need to be specified at a reference point in time, typically in terms of the instantaneous frequency




FIG. 5. Unlike the reference orbital phase ϕ_{ref} (top), the coordinate $\hat{\phi}_{\text{ref}}$ is well-measured (bottom), since it uniquely determines the observable phase of arrival at the reference detector φ_{k_0} . There remains, however, a discrete degeneracy between solutions separated by π that we treat in Sec. IV.

f_{ref} of the quadrupole radiation. Such parameters include orbital phase and, for precessing binaries, spin components, orbital inclination, and polarization. These are best constrained near frequencies where the signal is prominent; specifying them at a far-removed f_{ref} may introduce large correlations with other source parameters [28,29]. This effect is most important for the orbital phase, which in turn couples to the in-plane spin azimuths if they are defined relative to the orbital separation vector (we revisit this in Secs. III G and V B). We will choose the reference frequency to be $f_{\text{ref}} = \bar{f}_{k_0}^{\text{ML}}$; from Eq. (14), this is the frequency weighted by the squared signal-to-noise ratio, and hence it will naturally be within the detector band.

F. Aligned spin components

The effect of spins on the phase evolution of the inspiraling binary is largely dominated by the effective spin parameter

$$\chi_{\text{eff}} = \frac{\chi_{1z} + q\chi_{2z}}{1 + q}, \quad (24)$$

where χ_{1z}, χ_{2z} are the components of the constituent (dimensionless) spins parallel to the orbital angular momentum, and $q = m_2/m_1 \leq 1$ is the mass ratio. It is therefore convenient to use χ_{eff} directly as a spin coordinate. The Jacobian of the transformation to Cartesian spin components is nontrivial [30], however, note that the prior on spins is much less certain than that on extrinsic parameters. Our approach is to specify the prior in terms of χ_{eff} , which means we do not need to compute the Jacobian. We choose a uniform prior between ± 1 , which has the advantage that it does not vanish for any value of the observable χ_{eff} . In order to specify the two orbit-aligned

spin components, we also sample over their (typically) poorly measured difference, whose conditional prior we choose to be uniform within the range allowed by the cosmic censorship conjecture, $|\chi_{1,2}| < 1$ (see Appendix B). The prior induced on individual spin components is shown in [[31], Fig. 1].

For systems with well measured χ_{eff} , a common alternative description based on the spin magnitudes and tilts can suffer from correlations between these four variables.

[32,33] proposed a different parametrization which also removes the correlation of the aligned spins with the masses. We intend to implement this in the future.

G. In-plane spin components

In Sec. III F we have argued that \mathbf{L} defines one preferred axis to parametrize the spins, as their aligned components affect the evolution of the orbital phase. Describing the spin azimuths requires us to define a second axis, for which two choices are common in the literature: the orbital separation vector at f_{ref} , or the direction of propagation $\hat{\mathbf{N}} \equiv -\hat{\mathbf{n}}$. In line with [28], here we will advocate the latter choice.

Misaligned spins cause the orbital angular momentum \mathbf{L} , as well as the spins $\mathbf{S}_1, \mathbf{S}_2$, to precess about the total angular momentum \mathbf{J} , whose direction is stable [34,35]. This causes the inclination of the orbit seen from Earth, determined by $\hat{\mathbf{L}} \cdot \hat{\mathbf{N}}$, to continuously change. As a consequence, in addition to the usual cycles at twice the orbital frequency, misaligned-spin waveforms exhibit amplitude modulations at the slower precession rate. This separation of timescales means that the precession dynamics can be described using orbit-averaged equations, thereby decoupling it from the orbital phase.

Thus, the size and peak frequencies of the amplitude modulations are governed by the orbital and spin angular momenta at f_{ref} relative to $\hat{\mathbf{N}}$, as these determine the evolution of the inclination. On the other hand, per Eq. (12) the phase φ_k of the ‘‘carrier’’ wave is controlled by an approximately degenerate combination of $\phi_{\text{ref}}, \psi, \iota, \hat{\mathbf{n}}, t_c$. The practical consequence of this degeneracy is that ϕ_{ref} , which defines the azimuth about \mathbf{L} between the orbital separation and the direction of propagation, is poorly measured. In other words, a change in ϕ_{ref} can be compensated by adjusting e.g., ψ, d_L to keep $\hat{\phi}_{\text{ref}}, \hat{d}$ fixed, but only if the spins are held constant relative to $\hat{\mathbf{N}}$. If, instead, the spins are rigidly rotated with the binary, the peak frequencies of the observed amplitude modulations would shift in a way that cannot be compensated by adjusting other parameters. Thus, using the orbital separation to define the origin of spin azimuths would introduce in the parametrization a spurious coupling between the observable precession and orbital cycles. Following [28], we will use the angles $\theta_{JN}, \phi_{JL}, \phi_{12}$ to describe the inclination of the orbit and the spin azimuths. θ_{JN} and ϕ_{JL} are illustrated in Fig. 6: they define the direction of




FIG. 6. Angles describing relative orientation between angular momenta and the direction of wave propagation [28]. In terms of the observable precession cycles, changing $\hat{\mathbf{N}} \mapsto -\hat{\mathbf{N}}$ (i.e. inverting $\cos \theta_{JN}$ and adding π to ϕ_{JL}) is a symmetry that can cause a discrete degeneracy. We return to this point in Sec. IV A.

wave propagation $\hat{\mathbf{N}}$ given the total and orbital angular momenta \mathbf{J}, \mathbf{L} and irrespective of the orbital separation vector. ϕ_{12} is defined as the difference between the primary and secondary spin azimuths about \mathbf{L} . ϕ_{JL} posteriors are often bimodal, and for this reason we will eventually replace ϕ_{JL} by a modified spin azimuth $\hat{\phi}_{JL}$ in Sec. IV A.

Our parametrization of the spins is more akin to cylindrical than spherical coordinates, in the sense that we favor the orbit-aligned spin components over the zenithal angles. The description of spins is completed by specifying the in-plane spin magnitudes $\chi_1^\perp, \chi_2^\perp$.

IV. MITIGATING MULTIMODALITY

Oftentimes, the posterior distribution of gravitational wave source parameters is multimodal. Such distributions are more challenging for stochastic samplers, which might occasionally misestimate the relative weights of the modes or miss some of them altogether.

In Sec. IV A we identify four approximate discrete symmetries that are responsible for frequent multimodality in the orbital phase, polarization, inclination, and sky-location parameters. In Sec. IV B we introduce ‘‘folding’’; a simple algorithm to exploit the knowledge of the underlying approximate symmetries in order to robustly sample this type of multimodal probability distribution. Depending on the source properties, the detector network and the signal-to-noise ratio, this procedure can generically reduce the number of disjoint modes by a factor of up to eight.

A. Approximate discrete symmetries

In Sec. III D we already encountered an approximate discrete symmetry

$$\hat{\phi}_{\text{ref}} \mapsto \hat{\phi}_{\text{ref}} + \pi, \quad (25)$$

(Fig. 5) which is exact for waveforms with only even values of m , in particular the dominant mode $(\ell, |m|) = (2, 2)$.

A similar symmetry exists for the polarization

$$\psi \mapsto \psi + \frac{\pi}{2} \quad (26)$$

at constant $\hat{\phi}_{\text{ref}}$, which [per Eq. (22)] entails a simultaneous change by $\pi/2$ in ϕ_{ref} . This symmetry arises because, under the transformation (26), the antenna coefficients F_k^+, F_k^\times change sign [20]; meanwhile, the $\pi/2$ rotation of ϕ_{ref} inverts the sign of the waveforms $h_{\ell m}^+, h_{\ell m}^\times$ for modes with $|m| = 2$. These two sign flips cancel, leaving the measurable responses $h_k = F_k^+ h^+ + F_k^\times h^\times$ invariant. The symmetry becomes exact for waveforms with $|m| = 2$. Usually, this discrete degeneracy does not produce disjoint modes in the posterior, because the uncertainty in ψ is large enough that the two solutions remain connected (as an extreme example, ψ is a dummy parameter for face-on, aligned-spin waveforms). Still, ψ posteriors generally do exhibit this symmetry.

We identify two other approximate discrete symmetries. In Secs. III A–III D we derived coordinates $\hat{d}, \hat{\phi}_{\text{ref}}, t_{k_0}, \theta_{\text{net}}$, which determine the amplitude, phase and time of arrival at the leading detector, and arrival time at the second detector. These variables typically capture most of the information on extrinsic parameters from the likelihood, and consequently tend to be tightly constrained. Conversely, the remaining extrinsic parameters $\theta_{JN}, \phi_{\text{net}}, \psi$ have a smaller effect on the likelihood. As a result, they can exhibit large degeneracies—some of them discrete, leading to multiple modes. Indeed, due to the geometry of the Hanford-Livingston network, for these detectors both the prior and likelihood are approximately symmetric under either of the following discrete transformations (at fixed $\hat{d}, \hat{\phi}_{\text{ref}}, t_{k_0}, \theta_{\text{net}}$),

$$\phi_{\text{net}} \mapsto -\phi_{\text{net}}, \quad (27a)$$

$$(\phi_{\text{net}}, \cos \theta_{JN}, \phi_{JL}) \mapsto (\phi_{\text{net}} + \pi, -\cos \theta_{JN}, \phi_{JL} + \pi). \quad (27b)$$

The manifestation of these symmetries in the posterior is shown in Fig. 7. We can understand them intuitively as follows. By design, the Hanford and Livingston detectors are nearly coaligned (plus a $\pi/2$ rotation in the horizontal plane, which simply adds a phase difference of π ; see Fig. 3). For perfectly coaligned detectors, the amplitude and phase at the second detector are determined by those at the first; in this sense their measurement does not provide new constraints, and the likelihood is largely independent of $\phi_{\text{net}}, \iota, \psi$. These parameters are therefore significantly informed by the prior. The prior has nontrivial structure in




FIG. 7. Approximate discrete symmetries are frequently responsible for multiple modes in the distribution, especially for the Hanford-Livingston network due to its peculiar geometric configuration. The plot shows examples of posteriors for the cosine of the inclination and a coordinate $\hat{\phi}_{\text{net}}$ [Eq. (29)] that describes the line-of-sight azimuth along a ring of constant time delay. Additional detectors and higher modes can break these symmetries, as for GW190412 [36].

these variables because the distance, and thereby the observable volume, depends on $\iota, \phi_{\text{net}}, \psi$ at constant $\hat{d}, \theta_{\text{net}}$ (which are fixed by the likelihood). For example, the uniform prior in luminosity volume $\pi(d_L) \propto d_L^2$ is [by Eqs. (18) and (19)]

$$\pi(\hat{d} | \mathcal{M}, \iota, \theta_{\text{net}}, \phi_{\text{net}}, \psi) \propto \hat{d}^2 \mathcal{M}^{5/2} |R_{k_0}(\iota, \theta_{\text{net}}, \phi_{\text{net}}, \psi)|^3, \quad (28)$$

proportional to the cube of the absolute value of the reference detector response. This term has four peaks as a function of ι, ϕ_{net} near $\cos \iota = \pm 1$ [Eq. (9)] and $\phi_{\text{net}} = \pm \pi/2$ (Fig. 4); see also Fig. 13. In other words, corresponding to a source face-on or face-off, and above or beneath the detector (constrained to the time-delay ring). As a result, Hanford-Livingston posteriors typically exhibit four modes at these configurations. We can see that both transformations (27) map an overhead source to one underfoot, i.e., they send $\hat{n} \cdot \hat{y}_{\text{net}} \equiv n_y \mapsto -n_y$. The second transformation, Eq. (27b), also flips the inclination. All four configurations can be reached by applying combinations of these two transformations.

The transformation $\cos \theta_{JN}, \phi_{JL} \mapsto -\cos \theta_{JN}, \phi_{JL} + \pi$ in Eq. (27b) corresponds to inverting the direction of propagation $\hat{N} \mapsto -\hat{N}$ in the frame of the binary. In other words, it can be interpreted as observing the system from the antipodal direction. To the extent that the source can be modeled as a precessing quadrupole, this simply reverses the handedness of the gravitational wave and cannot be distinguished with coaligned detectors. In reality, due to the curvature of Earth the detectors are not perfectly coaligned. As a result, the phase of arrival at the second detector is not completely determined by that at the first detector and its measurement has some constraining power. Whether it is

advanced or retarded depends on the handedness of the elliptically polarized wave and the relative orientation of the two detectors as viewed from the source (the “projected detector tensors” [37]). For example, with the system defined in Fig. 3, the projected Hanford detector appears rotated counter-clockwise relative to Livingston when seen from the $x_{\text{net}} > 0$ hemisphere, but clockwise from $x_{\text{net}} < 0$. Thus, the phase at Hanford would be advanced relative to Livingston for a right-polarized wave from $n_x < 0$, or left-polarized wave from $n_x > 0$ (and vice versa), making these solutions degenerate. Thus, n_x has to be inverted simultaneously for the transformation $\hat{N} \mapsto -\hat{N}$ to be a good symmetry of the likelihood. This is implemented by sending $\phi_{\text{net}} \mapsto \phi_{\text{net}} + \pi$ in Eq. (27b).

Incidentally, since the arrival phase difference also involves the time delay, this effect induces an observable correlation between θ_{net} and ϕ_{net} . We provide a more quantitative treatment of all this in Appendix A.

In order to simplify the transformation (27b), we will define two 2π -periodic coordinates

$$\hat{\phi}_{\text{net}} := \begin{cases} \phi_{\text{net}} & \text{if } \cos \theta_{JN} < 0 \\ \phi_{\text{net}} + \pi & \text{else} \end{cases} \quad (29)$$

$$\hat{\phi}_{JL} := \begin{cases} \phi_{JL} & \text{if } \cos \theta_{JN} < 0 \\ \phi_{JL} + \pi & \text{else} \end{cases} \quad (30)$$

to replace ϕ_{net} and ϕ_{JL} in the characterization of the sky location and spin azimuths, respectively. Both transformations have unit Jacobian. With these coordinates, the approximate symmetries in Eq. (27) become one-parameter reflections,

$$\hat{\phi}_{\text{net}} \mapsto -\hat{\phi}_{\text{net}}, \quad (31a)$$

$$\cos \theta_{JN} \mapsto -\cos \theta_{JN}. \quad (31b)$$

$\hat{\phi}_{JL}$ has a simple interpretation as the azimuth about \mathbf{J} between \mathbf{L} and the unsigned direction of propagation, so that $\hat{N} \mapsto -\hat{N}$ leaves $\hat{\phi}_{JL}$ invariant. In Sec. VB we will show that this azimuth can be remarkably well-measured.

The Virgo detector does not have such special alignment. If a signal is loud in Virgo and at least one LIGO detector, the symmetries in Eq. (31) are typically broken and the posterior distribution in these variables is unimodal.

B. Folding algorithm

Having identified the approximate discrete symmetries responsible for multimodality, we now introduce “folding”, an algorithm to exploit this structure when sampling. The basic idea is illustrated in Fig. 8, beginning with sampling from the probability distribution marginalized over the discrete approximate-symmetry transformations, and then

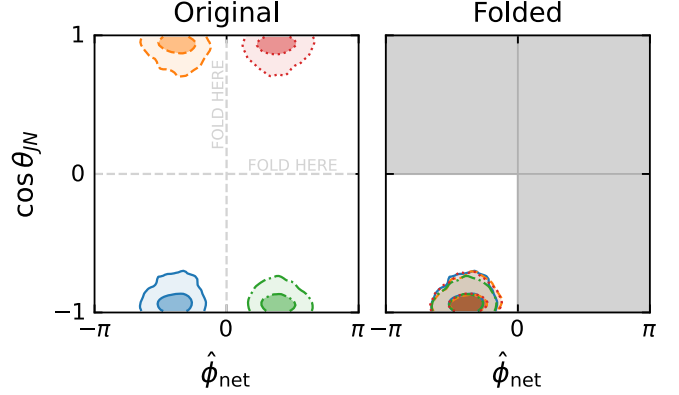


FIG. 8. Folding algorithm. Left: the posterior has multiple modes (shown in different colors) due to approximate symmetries that are known in advance. Right: we can “fold” the distribution (sum its appropriately transformed modes) to make it unimodal. We sample the folded distribution and reconstruct the original in postprocessing.

sampling over the set of transformations in postprocessing to undo the marginalization.

The four approximate symmetry transformations in Eqs. (25), (26), (31a), and (31b) allow us to divide the phase space into $2^4 = 16$ sectors in which the posterior has similar behavior. This is illustrated in the left panel of Fig. 8 for the four quadrants of $\cos \theta_{JN}$, $\hat{\phi}_{\text{net}}$ space, which have similar solutions by virtue of the approximate symmetries in Eq. (31) (the other two folded dimensions $\hat{\phi}_{\text{ref}}$, ψ are not shown). We pick one of these sectors, as highlighted in the right panel, which we call the “folded space”. The full space can be recovered from the folded space via 2^4 discrete mappings $\{\sigma_1, \dots, \sigma_{16}\}$ that either do or do not apply each of the transformations (25), (26), (31a), and (31b), respectively. If the symmetries were perfect, it would suffice to draw samples on the folded space, then distribute them evenly onto all sectors using these mappings. The symmetries in Sec. IVA are only approximate, since the detectors are not perfectly aligned, and general waveform models have additional effects like orbital precession and higher harmonics that are not modeled by Eq. (7). Hence we generalize this idea to relax the requirement of perfect symmetry, as follows:

- (1) Define the folded distribution

$$\tilde{P}(\tilde{\theta}) = \sum_{i=1}^{2^N} P(\sigma_i(\tilde{\theta})), \quad (32)$$

where $\tilde{\theta}$ is a set of parameters in the folded space, $N = 4$ is the number of folded parameters, $\{\sigma_i\}$ are the 2^N discrete mappings and P is the posterior distribution in the full space. The number of modes of \tilde{P} can be smaller than that of P , as shown in the right panel of Fig. 8, by a factor of up to 2^N .

- (2) Draw “folded” samples from this simpler distribution, $\{\tilde{\theta}^j\} \sim \tilde{P}$, using a traditional sampler.
- (3) Assign each folded sample to a sector: for each $\tilde{\theta}^j$, draw θ^j from $\{\sigma_1(\tilde{\theta}^j), \sigma_2(\tilde{\theta}^j), \dots\}$ with relative probabilities $\{P(\sigma_1(\tilde{\theta}^j)), P(\sigma_2(\tilde{\theta}^j)), \dots\}$. The set $\{\theta^j\}$ is distributed according to $P(\theta)$.

At first glance, Eq. (32) suggests that each evaluation of \tilde{P} requires 2^N evaluations of P and would therefore increase the computational cost by that factor. However, since the folded parameters are extrinsic, the expensive computation of the waveform can be reused. In the case at hand, changing the polarization and sky location requires recomputing antenna factors F_+ , F_\times and time delays, and changing the phase requires recomputing spherical harmonic phases $e^{im\phi_{\text{ref}}}$. Even in a general case, all 2^N evaluation points are known simultaneously, facilitating vectorization and parallelization. These considerations make the cost of computing \tilde{P} similar to that of P . Likewise, generating θ^j from $\tilde{\theta}^j$ in step 3 above does not require additional computations of P , because these values can be stored along with $\tilde{\theta}^j$ during the sampling process, and the assignment itself is computationally very cheap. In practice, the number of evaluations required for convergence is much smaller for distributions with less modes, making this method advantageous for both robustness and efficiency.

We emphasize that the folding procedure is not an approximation; it still gives the correct answer if the distribution P is not symmetric at all (in that case, it just does not provide any advantage).

In Sec. III D we had mentioned that the $\hat{\phi}_{\text{ref}}$ posterior is bimodal and discontinuous. In contrast, per Eqs. (22), (25), and (32) the folded posterior is unimodal and continuous, because both sides of the branch cut are summed together. In a similar way, the discontinuities introduced in Eqs. (29) and (30) are absent in the folded posterior since they happen at the fold $\cos\theta_{JN} = 0$.

Finally, one might worry that Fig. 8 merely shows that the marginalized 2-dimensional posterior is approximately symmetric, while the folding algorithm is efficient only if the full 15-dimensional posterior is symmetric. It could be possible that, although this projection looks symmetric, the mappings we proposed were incorrect descriptions of the symmetry in the full space. In Fig. 9 we show histograms of the unfolding probabilities $p_i \propto P(\sigma_i(\tilde{\theta}))$ used in step 3, which have unit sum by construction. In the limit that the transformations σ_i were perfect symmetries of P , these probabilities would be $1/16$. Conversely, if they were not good symmetries, these probabilities would be very nearly 0 or 1. We find that they are near $1/16$, confirming that the transformations we identified are indeed approximate symmetries of the full space. Whether all or some of the modes are present for any particular event depends on the source parameters, the network configuration and the signal-to-noise ratio.




FIG. 9. Distribution of probabilities p_i with which folded samples are assigned to the 16 sectors during step 3, for GW151012. All peak near $1/16$, demonstrating that the transformations (25), (26), (31a), and (31b) are approximate symmetries of the full 15-dimensional space.

An alternative to the folding algorithm, which is similar in spirit, is to specify custom jump proposals that follow the approximate symmetry transformations and effectively “connect” the modes in the sampling process. This approach has been used to mitigate both degeneracy and multimodality [38–40], albeit for a different set of symmetries [that includes the transformation in Eq. (25)]. The choice of sampling proposal can significantly impact the efficiency of Monte Carlo methods, see e.g., [41,42] and references therein.

V. RESULTS

We implement the coordinate transformations described in Sec. III and the folding algorithm introduced in Sec. IV in a PYTHON package, which we make publicly available at <https://github.com/jroulet/cogwheel>. Apart from those improvements, the `cogwheel` code utilizes the relative-binning algorithm [43–45] with higher-order modes [46] to accelerate likelihood evaluations. It interfaces with third-party routines for downloading public data (GWOSC [47], GWpy [48]), generating waveforms (lalsuite [49]) and sampling distributions (PyMultiNest [50,51], dynesty [52]).

In this section we summarize the improvements brought by the coordinate transformations and folding technique in terms of efficiency and robustness of the inference. Throughout, we model waveforms using IMRPhenomXPHM with next-to-next-to-leading-order post-Newtonian precession [17], and we sample distributions with PyMultiNest [50,51].

Figure 10 summarizes how our choice of coordinates, combined with the folding method, naturally describes the distribution of extrinsic parameters. The uncertainties achieved in the one-dimensional marginal posteriors for the parameters controlling distance, arrival phase and arrival times are in much closer agreement with the expectations from Eqs. (15)–(17) and (21), shown by

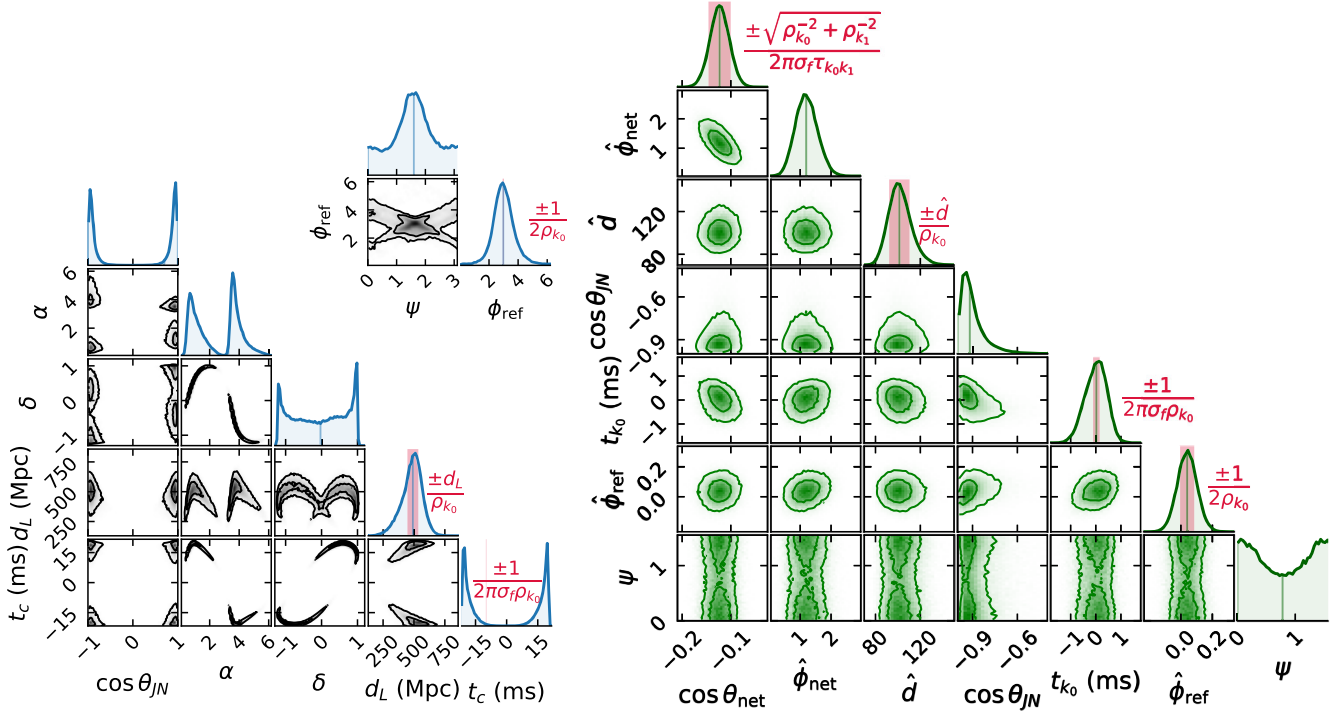


FIG. 10. Marginalized posterior for the extrinsic parameters of GW151226 in terms of standard coordinates (left) or the folded coordinates advocated in this work (right), see Secs. III and IV. The complex structures on the left are largely absent on the right, making the latter system better suited for sampling. Crimson bars show the standard deviation expected given the signal-to-noise ratio [Eqs. (15)–(17), (21)], discrepancies with the realized uncertainties are due to correlations with the marginalized parameters.

crimson bars. That being said, Eq. (17) significantly underestimates the uncertainty of the arrival time t_{k_0} . This can be traced back to correlations with intrinsic parameters, most significantly chirp mass and in-plane spin magnitude (not shown). The definition of arrival time for waveforms with different intrinsic parameters is somewhat arbitrary. The convention adopted by the LIGO Algorithm Library is to define the reference time when the amplitude of the strain is maximal [53]. In the IMRPhenomX family of waveform models this is implemented through parametric fits [54] that are accurate to ~ 1 ms [55], which is consistent with the precision of the arrival time measurement achieved in Fig. 10.

A. Performance of stochastic sampling

We compare the performance of parameter estimation runs using the coordinate system presented in Sec. III (coupled to the folding algorithm of Sec. IV B) against an “unoptimized” system that uses $(\cos \theta_{\text{net}}, \phi_{\text{net}}, t_{k_0}, \cos \theta_{JN}, \phi_{JL}, \phi_{12}, \cos \theta_1, \cos \theta_2, \chi_1, \chi_2, d_L, \phi_{\text{ref}}, \psi, \mathcal{M}, \ln q)$ as sampling coordinates and no folding. This comparison assesses the impact of the methods new to this work within the `cogwheel` software, note that other state-of-the-art codes for parameter inference [38–40, 56–58] use various other parametrizations and optimizations (marginalization, custom jump proposals, parallelization, etc.). We adopt an isotropic, uniform-in-spin-magnitudes prior in both cases, and so for ease of

implementation in the optimized case we use the cumulative of the prior on aligned spin components as sampling coordinate instead of the effective spin (cf. Sec. III F). In both examples we use `PyMultiNest` and vary the number of live points in factors of two between 512 and 16384 (this is an internal feature of nested samplers; a higher number of live points typically achieves a better coverage of parameter space and produces more independent samples at a proportionally higher computational cost). In all cases, we use a single computing core and identical environments to run each configuration, thus, differences in runtime are due to the algorithm. We use GW151226 as a test case.

We find that, despite nested sampling being generally well suited for multimodal problems, and despite using numerous live points, the sampler fails to find some modes of the posterior when folding is not used. This is illustrated in the top panel of Fig. 11; of the unoptimized runs, only the one with 16384 live points found all modes, and even in this case one of the modes was undersampled. Conversely, all runs that used our coordinate system and folding successfully found all the modes. To quantify the error this induces, we compare each posterior distribution P to a reference answer P_{ref} , chosen as that with the largest number of live points and using our coordinate system and folding algorithm. Our measure of the error is the Jensen-Shannon divergence between each marginal $\theta_{JN}, \phi_{\text{net}}$ posterior and the reference distribution,

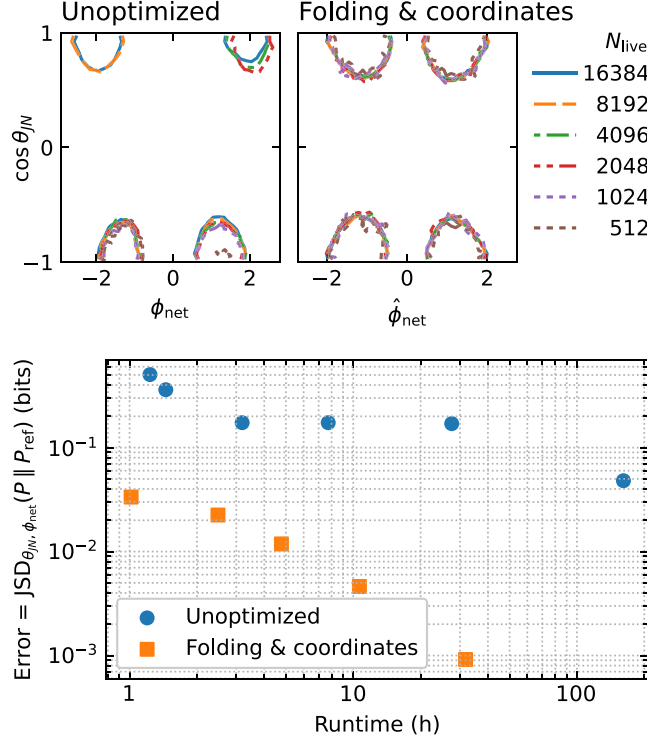


FIG. 11. Improvement brought by our methods in terms of increased robustness to multimodality. We compare runs that do not use the methods new to this work (“unoptimized”) to runs that do (“folding and coordinates”). Within each setting, runs differ in number of live points. *Top*: Without folding, the sampler fails to find all modes of the posterior even when using a large number of live points. *Bottom*: Error achieved in the posterior distribution versus runtime, as the number of live points is iteratively doubled. We use the Jensen-Shannon divergence with respect to the highest-resolution “folding and coordinates” run as our error measure, see Eq. (33) (the reference run is excluded from the figure). The unoptimized runs have a higher error due to missing or undersampled modes.

$$\text{JSD}_{\theta_{JN}, \phi_{\text{net}}}(P \parallel P_{\text{ref}}) := \frac{1}{2} \iint d\theta_{JN} d\phi_{\text{net}} \left(P \log_2 \frac{P}{M} + P_{\text{ref}} \log_2 \frac{P_{\text{ref}}}{M} \right) \quad (33)$$

with $M = (P + P_{\text{ref}})/2$. We singled out θ_{JN} , ϕ_{net} because this projection of the unoptimized posteriors is most clearly missing modes. To evaluate Eq. (33), we obtain $P(\theta_{JN}, \phi_{\text{net}})$, $P_{\text{ref}}(\theta_{JN}, \phi_{\text{net}})$ from the posterior samples with a kernel density estimation and perform a double quadrature, using `scipy` implementations [59]. We show this in the bottom panel of Fig. 11: we see more than an order of magnitude improvement in terms of this error measure when we use our coordinate system and folding.

As expected, the runtime is approximately proportional to the number of live points, however, the unoptimized runs show larger fluctuations. We attribute these to the stochastic nature of the sampling process, likely amplified by the fact




FIG. 12. Posterior distributions of modified spin azimuth $\hat{\phi}_{JL}$ [Eq. (30)] for the 15 events with best constraints, under a prior isotropic in spins. For aligned-spin configurations, $\hat{\phi}_{JL}$ becomes a dummy parameter, and therefore events with well-measured $\hat{\phi}_{JL}$ have spins misaligned with the orbit. Posteriors have been artificially centered by subtracting their circular means, which have no astrophysical importance.

the distribution is multimodal: the runtime can be significantly affected by whether, and when, a new mode is found by the sampler, especially when using few live points, for which the relative fluctuations are found to be largest.

B. Measurability of spin azimuth

From the astrophysics standpoint, the most interesting result of this work is that the modified spin azimuth $\hat{\phi}_{JL}$ can be measured remarkably well. Figure 12 shows its posterior distribution for the 15 events in the GWTC-3 and Institute for Advanced Study’s catalogs with tightest bounds. We see that it is not uncommon for this parameter to be significantly constrained. Recall that $\hat{\phi}_{JL}$ describes the azimuth about \mathbf{J} between the orbital angular momentum \mathbf{L} and the location of the detector [Eq. (30), Fig. 6]. Due to isotropy, all possible directions $\hat{\mathbf{N}}$ towards the detector are equivalent *a priori*, and thus the prior on $\hat{\phi}_{JL}$ is uniform. Moreover, for aligned-spin configurations we have $\mathbf{L} \parallel \mathbf{J}$, so $\hat{\phi}_{JL}$ has no effect on the waveform and therefore the likelihood is independent of $\hat{\phi}_{JL}$. It follows that any departure from a flat posterior originates from the presence of misaligned spins in the waveform model. Furthermore, if any particular value of $\hat{\phi}_{JL}$ can be confidently ruled out for an event, then the source is inconsistent with having aligned spins.

This can have far-reaching consequences for the astrophysical interpretation of binary mergers. The degree of spin-orbit alignment constrains the formation history of the system—in particular, whether it likely formed in isolation or in a dense environment [60]—and can inform astrophysical processes such as tidal interactions and supernova

kicks [61]. For this reason, characterizing the imprints of precession is a topic of intensive research. Indeed, beyond the parametrization in terms of θ_{JN} , ϕ_{JL} , ϕ_{12} on which ours is built [28], numerous other descriptions have been studied, both physically and observationally motivated; in terms of an effective precession spin χ_p [62] or modifications thereof [63,64], a precessing signal-to-noise ratio ρ_p [65], a taxonomy of phenomenological parameters [66], or the spin azimuths [29].

The practical difference between ϕ_{JL} and $\hat{\phi}_{JL}$ is that ϕ_{JL} posteriors typically have two modes separated by π , with opposite inclinations. The reason for this bimodality is an approximate discrete symmetry corresponding to observing the same source from the antipodal direction, as discussed in Sec. IV A. As a result, the marginal posterior on ϕ_{JL} is not particularly well constrained. We solve this by applying a shift of π depending on the sign of $\cos \theta_{JN}$ per Eq. (30).

Recently, [29] found hints of precession in GWTC-2 by studying the azimuths ϕ_1 , ϕ_2 about \mathbf{L} between each individual spin and the orbital separation vector, at a reference time $t_{\text{ref}} = -100GM/c^3$ before merger. The results in Fig. 12 are similar in spirit because, following an argument analogous to the one opening this section, a tight measurement of ϕ_1 or ϕ_2 would also rule out aligned spins. In fact, we find that significantly tighter constraints can be placed on $\hat{\phi}_{JL}$ than on ϕ_1 or ϕ_2 . This is for two reasons; first, the aforementioned symmetry leads to a similar bimodality correlated with the source inclination, and second, ϕ_1 and ϕ_2 define the spin azimuths relative to the poorly measured orbital phase ϕ_{ref} . Indeed, [29] report that varying the reference time (or equivalently, frequency) at which the source configuration is specified sensitively affects the constraints on ϕ_1 , ϕ_2 . Interestingly, we find that the width of the ϕ_{ref} posterior depends on f_{ref} , but that of $\hat{\phi}_{JL}$ does not, at least in the range $50 \text{ Hz} < f_{\text{ref}} < 150 \text{ Hz}$. This indicates that, at least in some cases, the uncertainty in ϕ_{ref} is the limiting factor in the measurement precision for the angle between spin and orbital separation. We interpret this observation as follows. First, while ϕ_1 , ϕ_2 evolve at the orbital timescale, the angle $\hat{\phi}_{JL}$ evolves at the slower precession timescale. Moreover, the precession frequency is well-constrained independently of directly observable precession effects, because intrinsic parameters are measured from the evolution of the orbital phase (note that to lowest post-Newtonian order, the precession frequency just depends on the masses and the orbital frequency, not the spins [28]). Since there are few precession cycles and their frequency is determined, their phase can be specified through $\hat{\phi}_{JL}$ similarly well over a broad range of reference frequencies. This said, we observe a degradation in its measurement if we adopt a reference frequency of 20 Hz. The azimuth difference ϕ_{12} also evolves on the precession timescale, however it has a subtler effect on the waveform and, in agreement with previous work [29,67], we find it is not as well-measured.

Reassuringly, the set of events with best measured $\hat{\phi}_{JL}$ contains those for which precession signatures have been reported before, namely GW200129_065458 [11,68], GW190412 [36,69,70], GW151226 [71], GW170818 [29], and GW190521 [72,73], as shown in Fig. 12. That being said, at this point we are unable to quantify the significance of the measurability of $\hat{\phi}_{JL}$ as an indicator of precession. This preliminary result motivates future work, in which we will define a quantitative statistic involving $\hat{\phi}_{JL}$ and calibrate its significance using synthetic data.

VI. CONCLUSIONS

We have introduced a coordinate system optimized for characterization of compact binary mergers observed through gravitational waves. It removes commonly encountered degeneracies and multimodality, and the transformation to standard coordinates has a simple Jacobian determinant and an explicit inverse. These coordinates improve the robustness and efficiency of parameter estimation algorithms, and build intuition about gravitational wave measurements.

In order to remove degeneracy, the coordinates are designed to separately control the main observable features of the signal. For the extrinsic parameters, these are the amplitude, phase and time of arrival at the reference (loudest) detector, and the time delay to the second-loudest detector. For the spins, we single out the effective spin parameter and the total spin azimuth, which affect the orbital evolution and the peak frequencies of the amplitude modulations induced by precession, respectively. We reemphasize previous realizations that the reference frequency at which the configuration of the binary is specified should be inside the sensitive band of the detectors, and the spin orientations should be defined relative to the direction of wave propagation rather than the orbital separation.

Strikingly, the azimuth $\hat{\phi}_{JL}$ between the total spin and the unsigned direction of wave propagation is well-measured in several examples (Fig. 12), establishing a novel observable signature of precession. We anticipate that this parametrization will improve our understanding of spin-orbit misalignment in nature. In future work we will define a statistic based on this parameter to quantify the significance of spin misalignment.

We identified four approximate symmetries as the leading cause of multimodality in posterior distributions, involving the orbital phase, polarization, sky location, and inclination. Depending on the signal-to-noise ratio and the network configuration, these can lead to up to 2^4 modes in the posterior. The last two of these symmetries are particularly good for the Hanford-Livingston network due to its peculiar geometric configuration, and can lead to four degenerate solutions roughly corresponding to a source overhead or underfoot, and face-on or face-off.

We devised “folding”, an algorithm that turns a distribution with this type of multimodality into an equivalent

unimodal problem by marginalizing over the discrete degeneracies. To facilitate its application, we adapted our coordinates so that each of the four independent approximate-symmetry transformations is achieved by a one-parameter shift or reflection.

Using these algorithms, we were able to achieve robust parameter inference while keeping the number of live points, and thereby the computational cost, low (Fig. 11). We make publicly available a parameter estimation code that implements these features at <https://github.com/jroulet/cogwheel>.

These methods greatly simplify the distribution, to the point where a parametric approximation to the full distribution can be made based on the signal-to-noise ratio and a few inputs specifying the location of the peak. Beyond the applications shown here, we expect that this will have other uses in gravitational wave data analysis. For example, evidence integrals have application in search [13,74] and parameter estimation [75], and knowledge of the distribution can be used to design efficient integration schemes, e.g., through variance reduction methods. As another example, a machine-learning approach based on normalizing flows was recently demonstrated to be accurate and fast at estimating gravitational wave source parameters [76]. In that approach, the posterior distribution is described in terms of a system of coordinates in which it has a standard form (e.g., Gaussian). This change of coordinates is found automatically and contains the complexity of the problem. The techniques introduced in this paper can be regarded as an analytical approximation to a normalizing flow. This suggests that applying a normalizing flow on the space of coordinates we developed here (or similar) might require a simpler neural network architecture and reduce the training cost.

ACKNOWLEDGMENTS

We thank Geraint Pratten for help with waveform models and Will Farr for comments on the paper. This research has made use of data or software obtained from the Gravitational Wave Open Science Center (gw-openscience.org), a service of LIGO Laboratory, the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA. LIGO Laboratory and Advanced LIGO are funded by the United States National Science Foundation (NSF) as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by institutions from

Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain. The construction and operation of KAGRA are funded by Ministry of Education, Culture, Sports, Science and Technology (MEXT), and Japan Society for the Promotion of Science (JSPS), National Research Foundation (NRF) and Ministry of Science and ICT (MSIT) in Korea, Academia Sinica (AS) and the Ministry of Science and Technology (MoST) in Taiwan. J. R. is supported by Grant No. 216179 to the KITP from the Simons Foundation. S. O. acknowledges support as an NSF Graduate Research Fellow under Grant No. DGE-2039656. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. T. I. is supported in part by the Heising-Simons Foundation, the Simons Foundation, and National Science Foundation Grant No. NSF PHY-1748958 and Grant No. PHY-2110496. T. V. acknowledges support by the National Science Foundation under Grant No. 2012086. B. Z. is supported by a research grant from the Center for New Scientists at the Weizmann Institute of Science and a research grant from the Ruth and Herman Albert Scholarship Program for New Scientists. M. Z. is supported by the Canadian Institute for Advanced Research (CIFAR) program on Gravity and the Extreme Universe and the Simons Foundation Modern Inflationary Cosmology initiative.

APPENDIX A: APPROXIMATE SYMMETRIES, CONTINUED

In this appendix we provide a more quantitative discussion of the argument presented in Sec. IV A, regarding how the measurement of the Hanford-Livingston phase difference constrains the source azimuth along the time-delay ring and the discrete degeneracies associated.

From Eqs. (12) and (20), the observable phase difference between detectors k_0 and k_1 is

$$\begin{aligned} \varphi_{k_1} - \varphi_{k_0} = & \arg R_{k_1} - \arg R_{k_0} - 2\pi \tilde{f}_{k_1} \tau_{k_0 k_1} \cos \theta_{\text{net}} \\ & - 2\pi (\tilde{f}_{k_0} - \tilde{f}_{k_1}) t_{k_0}, \end{aligned} \quad (\text{A1})$$

where the last term can be neglected if the detectors have similar noise power spectrum shapes and thus \tilde{f}_k . The difference $\arg R_{k_1} - \arg R_{k_0}$ does depend on the sky location and inclination if the detectors are not coaligned, which provides a joint constraint on $\theta_{\text{net}}, \phi_{\text{net}}, \iota$. Likewise, the relative amplitude at the detectors

$$\frac{a_{k_1}}{a_{k_0}} = \left| \frac{R_{k_1}}{R_{k_0}} \right| \quad (\text{A2})$$

is also measurable. Figure 13 shows how these terms depend on $\hat{\phi}_{\text{net}}$ for various inclinations (holding fixed the values of our other coordinates, and for an aligned-spin configuration so $\iota = \theta_{JN}$). Inverting $\cos \iota$ at fixed ϕ_{net} inverts the sign of




FIG. 13. Response of the Hanford-Livingston network along a ring of constant time-delay, as a function of modified azimuth $\hat{\phi}_{\text{net}}$ [Eq. (29)] for various inclinations. *Top*: prior probability, see Eq. (28). *Center*: phase difference introduced by the relative orientations of the detectors. *Bottom*: relative amplitude response. In these plots, the transformation (31a) is a horizontal reflection and (31b) inverts the color scale. Both are approximate symmetries of all three quantities.

$\arg R_k$ and adds π to $\hat{\phi}_{\text{net}}$, by Eqs. (9) and (29). Therefore, in the second panel blue and red curves are related by a vertical reflection plus a horizontal shift. As a result of the geometry of the Hanford-Livingston network, the relative phase and amplitude of the detector responses are symmetric under the transformations (31) to a good approximation, especially near the configurations favored by the prior $\hat{\phi}_{\text{net}} \approx \pm\pi/2$, $\cos i \approx \pm 1$.

The dependence of $\arg R_L - \arg R_H$ on $\hat{\phi}_{\text{net}}$ seen in Fig. 13 is responsible for the correlation between $\cos \theta_{\text{net}}$

and $\hat{\phi}_{\text{net}}$ noticeable in Fig. 10; per Eq. (A1), in order to match the observed $\varphi_{k_1} - \varphi_{k_0}$, $\cos \theta_{\text{net}}$ needs to change with $\hat{\phi}_{\text{net}}$ to compensate the variation of $\arg R_L - \arg R_H$.

APPENDIX B: REFERENCE SHEET

In this appendix we provide a compact list of the coordinates defined throughout the main text, as implemented in the `cogwheel` package. The sequence of transformations to a standard system is given in Table I, together with their Jacobian determinants. Here, GMST refers to the Greenwich Mean Sidereal Time of the event, which determines the orientation of Earth, and k_0, k_1 index the two detectors with highest signal-to-noise ratio.

We use $\chi_{1x}^N, \chi_{1y}^N, \chi_{2x}^N, \chi_{2y}^N$ as “standard” parameters to describe the in-plane spins; we define these as the Cartesian spins in a frame where $\hat{z} \parallel L$ and \hat{N} is in the yz -plane. This is related to the “radiation frame” used by the LIGO algorithm library [49] (where $\hat{z} \parallel L$ and \hat{x} is parallel to the orbital separation vector) with a rotation by ϕ_{ref} around \hat{z} ,

$$\begin{pmatrix} \chi_{1x} & \chi_{2x} \\ \chi_{1y} & \chi_{2y} \end{pmatrix} = \begin{pmatrix} \cos \phi_{\text{ref}} & \sin \phi_{\text{ref}} \\ -\sin \phi_{\text{ref}} & \cos \phi_{\text{ref}} \end{pmatrix} \begin{pmatrix} \chi_{1x}^N & \chi_{2x}^N \\ \chi_{1y}^N & \chi_{2y}^N \end{pmatrix}. \quad (\text{B1})$$

Using this system has two advantageous properties: the transformations in Table I get more decoupled since these spins are independent of ϕ_{ref} , and the coprocessing-frame harmonic modes of a waveform transform under a change by ϕ_{ref} as $h_{\ell m}(f) \rightarrow h_{\ell m} e^{im\phi_{\text{ref}}}$ if these spin components are held constant [17,77] which is useful for reusing waveform computations, e.g., in folding.

We also introduced spin coordinates $C_{\text{diff}}, C_1^\perp, C_2^\perp$ which are the cumulatives of the prior on the aligned spin difference and the in-plane spin magnitudes, respectively. These have a uniform prior on $(0,1)$ by definition, and their relation to the physical spins depends on the choice of prior.

TABLE I. Modular sequence of transformations from the sampling coordinates we propose to a standard system. Each transformation involves few variables, is only conditioned on variables computed by the previous ones, and has a simple Jacobian determinant.

Sampled	Standard	Conditioned on		$ J $	Reference
		Variables	Constants		
$\mathcal{M}, \ln q$	m_1, m_2			$\mathcal{M} \cosh^{2/5}(\frac{1}{2} \ln q)$	
$\chi_{\text{eff}}, C_{\text{diff}}$	χ_{1z}, χ_{2z}	q		^a	(24), (B2)
$\cos \theta_{JN}, \hat{\phi}_{JL}, \phi_{12}, C_1^\perp, C_2^\perp$	$\iota, \chi_{1x}^N, \chi_{1y}^N, \chi_{2x}^N, \chi_{2y}^N$	$\chi_{1z}, \chi_{2z}, m_1, m_2$	f_{ref}	^b	[28], (30), (B4)
ψ	ψ			1	
$\cos \theta_{\text{net}}, \hat{\phi}_{\text{net}}$	$\alpha, \sin \delta$	θ_{JN}	GMST, k_0, k_1	1	Fig. 3, (29)
t_{k_0}	t_c	α, δ	GMST, k_0	1	(13)
$\hat{\phi}_{\text{ref}}$	ϕ_{ref}	$t_c, \alpha, \delta, \psi, \iota$	GMST, $k_0, \tilde{f}_{k_0}^{\text{ML}}, \varphi_{k_0}^{\text{ML}}$	1	(22)
\hat{d}	d_L	$\mathcal{M}, \alpha, \delta, \psi, \iota$	GMST, k_0	\hat{d}/d_L	(18)

^aUnnecessary if the prior is chosen in terms of sampled parameters, e.g. uniform.

^bUnnecessary since the prior is uniform on sampled parameters by isotropy and the definition of C_1^\perp, C_2^\perp .

Our default choice is uniform in the aligned spin difference conditional on χ_{eff} , q , in which case

$$\begin{aligned} C_{\text{diff}}(\chi_{1z}, \chi_{2z}, q) &:= \int_{\chi_{1z}^{\min}}^{\chi_{1z}} d\chi'_{1z} \pi(\chi'_{1z} | \chi_{\text{eff}}, q) \\ &= \frac{\chi_{1z} - \chi_{1z}^{\min}}{\chi_{1z}^{\max} - \chi_{1z}^{\min}}, \end{aligned} \quad (\text{B2})$$

where $\chi_{1z}^{\min}(\chi_{1z}, \chi_{2z}, q)$ is the minimum possible χ_{1z} consistent with the value of effective spin $\chi_{\text{eff}}(\chi_{1z}, \chi_{2z}, q)$ [given by Eq. (24)] subject to the Kerr bound $|\chi_{1z}| < 1$, and similarly $\chi_{1z}^{\max}(\chi_{1z}, \chi_{2z}, q)$ is the maximum possible value,

$$\begin{aligned} \chi_{1z}^{\min} &= \max(\chi_{1z} + q\chi_{2z} - q, -1), \\ \chi_{1z}^{\max} &= \min(\chi_{1z} + q\chi_{2z} + q, 1). \end{aligned} \quad (\text{B3})$$

For the in-plane spins, our prior choice is uniform in the disk given the aligned spin value, which yields

$$\begin{aligned} C_1^\perp(\chi_1^\perp, \chi_{1z}) &:= \int_0^{\chi_1^\perp} d\chi_1^{\perp'} \pi(\chi_1^{\perp'} | \chi_{1z}) \\ &= \frac{(\chi_1^\perp)^2}{1 - \chi_{1z}^2}, \end{aligned} \quad (\text{B4})$$

and similarly for the secondary in-plane spin.

-
- [1] J. Aasi *et al.*, Advanced LIGO, *Classical Quantum Gravity* **32**, 074001 (2015).
- [2] F. Acernese *et al.*, Advanced Virgo: A second-generation interferometric gravitational wave detector, *Classical Quantum Gravity* **32**, 024001 (2015).
- [3] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo During the First and Second Observing Runs, *Phys. Rev. X* **9**, 031040 (2019).
- [4] A. H. Nitz, C. Capano, A. B. Nielsen, S. Reyes, R. White, D. A. Brown, and B. Krishnan, 1-OGC: The first open gravitational-wave catalog of binary mergers from analysis of public Advanced LIGO data, *Astrophys. J.* **872**, 195 (2019).
- [5] B. Zackay, T. Venumadhav, L. Dai, J. Roulet, and M. Zaldarriaga, Highly spinning and aligned binary black hole merger in the Advanced LIGO first observing run, *Phys. Rev. D* **100**, 023007 (2019).
- [6] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, New binary black hole mergers in the second observing run of Advanced LIGO and Advanced Virgo, *Phys. Rev. D* **101**, 083030 (2020).
- [7] A. H. Nitz, T. Dent, G. S. Davies, S. Kumar, C. D. Capano, I. Harry, S. Mozzon, L. Nuttall, A. Lundgren, and M. Tápai, 2-OGC: Open gravitational-wave catalog of binary mergers from analysis of public Advanced LIGO and Virgo data, *Astrophys. J.* **891**, 123 (2020).
- [8] B. Zackay, L. Dai, T. Venumadhav, J. Roulet, and M. Zaldarriaga, Detecting gravitational waves with disparate detector responses: Two new binary black hole mergers, *Phys. Rev. D* **104**, 063030 (2021).
- [9] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run, *Phys. Rev. X* **11**, 021053 (2021).
- [10] R. Abbott *et al.* (The LIGO Scientific and the Virgo Collaborations), GWTC-2.1: Deep extended catalog of compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run, [arXiv:2108.01045](https://arxiv.org/abs/2108.01045).
- [11] R. Abbott *et al.* (The LIGO Scientific, Virgo and KAGRA Collaborations), GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run, [arXiv:2111.03606](https://arxiv.org/abs/2111.03606).
- [12] A. H. Nitz, C. D. Capano, S. Kumar, Y.-F. Wang, S. Kastha, M. Schäfer, R. Dhurkunde, and M. Cabero, 3-OGC: Catalog of gravitational waves from compact-binary mergers, *Astrophys. J.* **922**, 76 (2021).
- [13] S. Olsen, T. Venumadhav, J. Mushkin, J. Roulet, B. Zackay, and M. Zaldarriaga, New binary black hole mergers in the LIGO–Virgo O3a data, *Phys. Rev. D* **106**, 043009 (2022).
- [14] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, D. Gerosa, L. C. Stein, L. E. Kidder, and H. P. Pfeiffer, Surrogate models for precessing binary black hole simulations with unequal masses, *Phys. Rev. Res.* **1**, 033015 (2019).
- [15] S. Ossokine, A. Buonanno, S. Marsat, R. Cotesta, S. Babak, T. Dietrich, R. Haas, I. Hinder, H. P. Pfeiffer, M. Pürrer, C. J. Woodford, M. Boyle, L. E. Kidder, M. A. Scheel, and B. Szilágyi, Multipolar effective-one-body waveforms for precessing binary black holes: Construction and validation, *Phys. Rev. D* **102**, 044055 (2020).
- [16] S. Khan, F. Ohme, K. Chatziioannou, and M. Hannam, Including higher order multipoles in gravitational-wave models for precessing binary black holes, *Phys. Rev. D* **101**, 024056 (2020).
- [17] G. Pratten, C. García-Quirós, M. Colleoni, A. Ramos-Buades, H. Estellés, M. Mateu-Lucena, R. Jaume, M. Haney, D. Keitel, J. E. Thompson, and S. Husa, Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes, *Phys. Rev. D* **103**, 104056 (2021).
- [18] P. Jaranowski and A. Królak, Gravitational-wave data analysis. formalism and sample applications: The Gaussian case, *Living Rev. Relativity* **8**, 3 (2005).

- [19] C. Cutler and É. E. Flanagan, Gravitational waves from merging compact binaries: How accurately can one extract the binary’s parameters from the inspiral waveform?, *Phys. Rev. D* **49**, 2658 (1994).
- [20] J. T. Whelan, The geometry of gravitational wave detection, https://dcc.ligo.org/public/0106/T1300666/003/Whelan_geometry.pdf (2013).
- [21] S. Fairhurst, Triangulation of gravitational wave sources with a network of detectors, *New J. Phys.* **11**, 123006 (2009).
- [22] J. Roulet, L. Dai, T. Venumadhav, B. Zackay, and M. Zaldarriaga, Template bank for compact binary coalescence searches in gravitational wave data: A general geometric placement algorithm, *Phys. Rev. D* **99**, 123022 (2019).
- [23] P. R. Brady and S. Fairhurst, Interpreting the results of searches for gravitational waves from coalescing binaries, *Classical Quantum Gravity* **25**, 105002 (2008).
- [24] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence, *Phys. Rev. Lett.* **116**, 241103 (2016).
- [25] L. P. Singer and L. R. Price, Rapid Bayesian position reconstruction for gravitational-wave transients, *Phys. Rev. D* **93**, 024013 (2016).
- [26] I. M. Romero-Shaw *et al.*, Bayesian inference for compact binary coalescences with BILBY: Validation and application to the first LIGO–Virgo gravitational-wave transient catalogue, *Mon. Not. R. Astron. Soc.* **499**, 3295 (2020).
- [27] J. Veitch and A. Vecchio, Bayesian coherent analysis of inspiral gravitational wave signals with a detector network, *Phys. Rev. D* **81**, 062003 (2010).
- [28] B. Farr, E. Ochsner, W. M. Farr, and R. O’Shaughnessy, A more effective coordinate system for parameter estimation of precessing compact binaries from gravitational waves, *Phys. Rev. D* **90**, 024018 (2014).
- [29] V. Varma, M. Isi, S. Biscoveanu, W. M. Farr, and S. Vitale, Measuring binary black hole orbital-plane spin orientations, *Phys. Rev. D* **105**, 024045 (2022).
- [30] T. A. Callister, A thesaurus for common priors in gravitational-wave astronomy (2021), [arXiv:2104.09508](https://arxiv.org/abs/2104.09508).
- [31] Y. Huang, C.-J. Haster, J. Roulet, S. Vitale, A. Zimmerman, T. Venumadhav, B. Zackay, L. Dai, and M. Zaldarriaga, Source properties of the lowest signal-to-noise-ratio binary black hole detections, *Phys. Rev. D* **102**, 103024 (2020).
- [32] E. Lee, S. Morisaki, and H. Tagoshi, Mass-spin reparametrization for a rapid parameter estimation of inspiral gravitational-wave signals, *Phys. Rev. D* **105**, 124057 (2022).
- [33] S. Morisaki and V. Raymond, Rapid parameter estimation of gravitational waves from binary neutron star coalescence using focused reduced order quadrature, *Phys. Rev. D* **102**, 104020 (2020).
- [34] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, Spin-induced orbital precession and its modulation of the gravitational waveforms from merging binaries, *Phys. Rev. D* **49**, 6274 (1994).
- [35] L. E. Kidder, Coalescing binary systems of compact objects to (post)^{5/2}-Newtonian order. V. Spin effects, *Phys. Rev. D* **52**, 821 (1995).
- [36] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GW190412: Observation of a binary-black-hole coalescence with asymmetric masses, *Phys. Rev. D* **102**, 043015 (2020).
- [37] J. T. Whelan, Visualization of antenna pattern factors via projected detector tensors, <https://dcc.ligo.org/public/0068/T1100431/002/projectedTensor.pdf> (2012).
- [38] J. Veitch *et al.*, Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library, *Phys. Rev. D* **91**, 042003 (2015).
- [39] G. Ashton and C. Talbot, BILBY-MCMC: An MCMC sampler for gravitational-wave inference, *Mon. Not. R. Astron. Soc.* **507**, 2037 (2021).
- [40] N. J. Cornish and T. B. Littenberg, Bayeswave: Bayesian inference for gravitational wave bursts and instrument glitches, *Classical Quantum Gravity* **32**, 135012 (2015).
- [41] S. Sanjib, Markov Chain Monte Carlo methods for Bayesian data analysis in astronomy, *Annu. Rev. Astron. Astrophys.* **55**, 213 (2017).
- [42] D. W. Hogg and D. Foreman-Mackey, Data analysis recipes: Using Markov Chain Monte Carlo, *Astrophys. J. Suppl. Ser.* **236**, 11 (2018).
- [43] N. J. Cornish, Fast Fisher matrices and lazy likelihoods, [arXiv:1007.4820](https://arxiv.org/abs/1007.4820).
- [44] B. Zackay, L. Dai, and T. Venumadhav, Relative binning and fast likelihood evaluation for gravitational wave parameter estimation, [arXiv:1806.08792](https://arxiv.org/abs/1806.08792).
- [45] N. J. Cornish, Heterodyned likelihood for rapid gravitational wave parameter inference, *Phys. Rev. D* **104**, 104054 (2021).
- [46] N. Leslie, L. Dai, and G. Pratten, Mode-by-mode relative binning: Fast likelihood estimation for gravitational waveforms with spin-orbit precession and multiple harmonics, *Phys. Rev. D* **104**, 123030 (2021).
- [47] R. Abbott *et al.*, Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo, *SoftwareX* **13**, 100658 (2021).
- [48] D. M. Macleod, J. S. Areeda, S. B. Coughlin, T. J. Massinger, and A. L. Urban, GWpy: A PYTHON package for gravitational-wave astrophysics, *SoftwareX* **13**, 100657 (2021).
- [49] LIGO Scientific Collaboration, LIGO Algorithm Library (2018).
- [50] J. Buchner, A. Georgakakis, K. Nandra, L. Hsu, C. Rangel, M. Brightman, A. Merloni, M. Salvato, J. Donley, and D. Kocevski, X-ray spectral modelling of the AGN obscuring region in the CDFS: Bayesian model selection and catalogue, *Astron. Astrophys.* **564**, A125 (2014).
- [51] F. Feroz, M. P. Hobson, and M. Bridges, MultiNest: An efficient and robust Bayesian inference tool for cosmology and particle physics, *Mon. Not. R. Astron. Soc.* **398**, 1601 (2009).
- [52] J. S. Speagle, DYNESTY: A dynamic nested sampling package for estimating Bayesian posteriors and evidences, *Mon. Not. R. Astron. Soc.* **493**, 3132 (2020).
- [53] P. Schmidt, I. W. Harry, and H. P. Pfeiffer, Numerical relativity injection infrastructure (2017), [arXiv:1703.01076](https://arxiv.org/abs/1703.01076).
- [54] C. García-Quirós, M. Colleoni, S. Husa, H. Estellés, G. Pratten, A. Ramos-Buades, M. Mateu-Lucena, and R. Jaume, Multimode frequency-domain model for the

- gravitational wave signal from nonprecessing black-hole binaries, *Phys. Rev. D* **102**, 064002 (2020).
- [55] G. Pratten (private communication).
- [56] J. Lange, R. O’Shaughnessy, and M. Rizzo, Rapid and accurate parameter inference for coalescing, precessing compact binaries, 2018, [arXiv:1805.10457](https://arxiv.org/abs/1805.10457)
- [57] C. M. Biwer, C. D. Capano, S. De, M. Cabero, D. A. Brown, A. H. Nitz, and V. Raymond, PyCBC inference: A PYTHON-based parameter estimation toolkit for compact binary coalescence signals, *Publ. Astron. Soc. Pac.* **131**, 024503 (2019).
- [58] G. Ashton *et al.*, BILBY: A user-friendly bayesian inference library for gravitational-wave astronomy, *Astrophys. J. Suppl. Ser.* **241**, 27 (2019).
- [59] P. Virtanen *et al.*, SciPy 1.0: Fundamental algorithms for scientific computing in PYTHON, *Nat. Methods* **17**, 261 (2020).
- [60] B. Farr, D. E. Holz, and W. M. Farr, Using spin to understand the formation of LIGO and Virgo’s black holes, *Astrophys. J.* **854**, L9 (2018).
- [61] D. Gerosa, E. Berti, R. O’Shaughnessy, K. Belczynski, M. Kesden, D. Wysocki, and W. Gladysz, Spin orientations of merging black holes formed from the evolution of stellar binaries, *Phys. Rev. D* **98**, 084036 (2018).
- [62] P. Schmidt, F. Ohme, and M. Hannam, Towards models of gravitational waveforms from generic binaries: II. Modelling precession effects with a single effective precession parameter, *Phys. Rev. D* **91**, 024043 (2015).
- [63] D. Gerosa, M. Mould, D. Gangardt, P. Schmidt, G. Pratten, and L. M. Thomas, A generalized precession parameter χ_p to interpret gravitational-wave data, *Phys. Rev. D* **103**, 064067 (2021).
- [64] L. M. Thomas, P. Schmidt, and G. Pratten, New effective precession spin for modeling multimodal gravitational waveforms in the strong-field regime, *Phys. Rev. D* **103**, 083022 (2021).
- [65] S. Fairhurst, R. Green, M. Hannam, and C. Hoy, When will we observe binary black holes precessing?, *Phys. Rev. D* **102**, 041302 (2020).
- [66] D. Gangardt, N. Steinle, M. Kesden, D. Gerosa, and E. Stoikos, A taxonomy of black-hole binary spin precession and nutation, *Phys. Rev. D* **103**, 124026 (2021).
- [67] S. Biscoveanu, M. Isi, V. Varma, and S. Vitale, Measuring the spins of heavy binary black holes, *Phys. Rev. D* **104**, 103018 (2021).
- [68] M. Hannam *et al.*, General-relativistic precession in a black-hole binary, *Nature* **610**, 652 (2021).
- [69] C. Hoy, C. Mills, and S. Fairhurst, Evidence for subdominant multipole moments and precession in merging black-hole-binaries from GWTC-2.1, *Phys. Rev. D* **106**, 023019 (2022).
- [70] T. Islam, S. E. Field, C.-J. Haster, and R. Smith, Improved analysis of GW190412 with a precessing numerical relativity surrogate waveform model, *Phys. Rev. D* **103**, 104027 (2021).
- [71] H. S. Chia, S. Olsen, J. Roulet, L. Dai, T. Venumadhav, B. Zackay, and M. Zaldarriaga, Signs of higher multipoles and orbital precession in GW151226, *Phys. Rev. D* **106**, 024009 (2022).
- [72] R. Abbott *et al.*, Properties and astrophysical implications of the 150 M_\odot binary black hole merger GW190521, *Astrophys. J.* **900**, L13 (2020).
- [73] S. Olsen, J. Roulet, H. S. Chia, L. Dai, T. Venumadhav, B. Zackay, and M. Zaldarriaga, Mapping the likelihood of GW190521 with diverse mass and spin priors, *Phys. Rev. D* **104**, 083036 (2021).
- [74] R. Smith and E. Thrane, Optimal Search for an Astrophysical Gravitational-Wave Background, *Phys. Rev. X* **8**, 021019 (2018).
- [75] C. Pankow, P. Brady, E. Ochsner, and R. O’Shaughnessy, Novel scheme for rapid parallel parameter estimation of gravitational waves from compact binary coalescences, *Phys. Rev. D* **92**, 023002 (2015).
- [76] M. Dax, S. R. Green, J. Gair, J. H. Macke, A. Buonanno, and B. Schölkopf, Real-Time Gravitational Wave Science with Neural Posterior Estimation, *Phys. Rev. Lett.* **127**, 241103 (2021).
- [77] A. Bohé, M. Hannam, S. Husa, F. Ohme, M. Pürrer, and P. Schmidt, PhenomPv2—technical notes for the LAL implementation, https://dcc.ligo.org/public/0122/T1500602/004/PhenomPv2_technicalnotes.pdf (2016).