Bayesian inference for compact binary coalescences with BILBY: validation and application to the first LIGO–Virgo gravitational-wave transient catalogue

I. M. Romero-Shaw^{1,2,★}, C. Talbot^{1,2,3}, S. Biscoveanu^{1,4,5}, V. D’Emilio^{1,6}, G. Ashton^{1,2}, C. P. L. Berry^{1,7,8}, S. Coughlin^{1,7}, S. Galadage^{1,2}, C. Hoy^{1,6}, M. Hübner^{1,2}, K. S. Phukon^{1,9,10,11,12}, M. Pitkin^{1,13}, M. Rizzo^{1,7}, N. Sarin^{1,2}, R. Smith^{1,2}, S. Stevenson^{1,14,15}, A. Vajpeyi^{1,2}, M. Arène^{1,16}, K. Athar¹, S. Banagiri^{1,17}, N. Bose^{1,18}, M. Carney^{1,7}, K. Chatziioannou^{1,19}, J. A. Clark^{1,20}, M. Colleoni^{1,21}, R. Cotesta^{1,22}, B. Edelman^{1,23}, H. Estellés^{1,21}, C. García-Quirós^{1,21}, Abhirup Ghosh^{1,22}, R. Green^{1,6}, C.-J. Haster^{1,4,5}, S. Husa^{1,21}, D. Keitel^{1,21}, A. X. Kim^{1,7}, F. Hernandez-Vivanco^{1,2}, I. Magaña Hernandez^{1,24}, C. Karathanasis^{1,25}, P. D. Lasky^{1,2}, N. De Lillo^{1,8}, M. E. Lower^{1,14,15,26}, D. Macleod^{1,6}, M. Mateu-Lucena^{1,21}, A. Miller^{1,27,28,29}, M. Millhouse^{1,30,31}, S. Morisaki^{1,32}, S. H. Oh^{1,33}, S. Ossokine^{1,22}, E. Payne^{1,2}, J. Powell^{1,14,15}, G. Pratten^{1,34}, M. Pürrer^{1,22}, A. Ramos-Buades^{1,21}, V. Raymond^{1,6}, E. Thrane^{1,2}, J. Veitch^{1,8}, D. Williams^{1,8}, M. J. Williams^{1,8} and L. Xiao^{1,3}

Affiliations are listed at the end of the paper

Accepted 2020 September 11. Received 2020 September 8; in original form 2020 June 1

ABSTRACT

Gravitational waves provide a unique tool for observational astronomy. While the first LIGO–Virgo catalogue of gravitational-wave transients (GWTC-1) contains 11 signals from black hole and neutron star binaries, the number of observations is increasing rapidly as detector sensitivity improves. To extract information from the observed signals, it is imperative to have fast, flexible, and scalable inference techniques. In a previous paper, we introduced BILBY: a modular and user-friendly Bayesian inference library adapted to address the needs of gravitational-wave inference. In this work, we demonstrate that BILBY produces reliable results for simulated gravitational-wave signals from compact binary mergers, and verify that it accurately reproduces results reported for the 11 GWTC-1 signals. Additionally, we provide configuration and output files for all analyses to allow for easy reproduction, modification, and future use. This work establishes that BILBY is primed and ready to analyse the rapidly growing population of compact binary coalescence gravitational-wave signals.

Key words: gravitational waves – methods: data analysis – stars: neutron – stars: black holes – transients: black hole mergers – transients: neutron star mergers.

1 INTRODUCTION

Gravitational-wave astronomy presents a revolutionary opportunity to probe fundamental physics and astrophysics, ranging from the neutron star equation of state and stellar evolution to the expansion of the Universe. The first direct observations of gravitational-wave signals have been made by Advanced LIGO (Aasi et al. 2015) and Advanced Virgo (Acernese et al. 2015); their first gravitational-wave catalogue of transients (GWTC-1; Abbott et al. 2019f) contains 10 binary black hole coalescences and one binary neutron star coalescence. The third observing run may yield $\mathcal{O}(10^2)$ additional observations (Abbott et al. 2018b), with signals from a second binary neutron star merger (Abbott et al. 2020d), one merger of a black hole with a $2.6 M_{\odot}$ compact object, and an additional two binary black hole mergers (Abbott et al. 2020a, c) already confirmed.

Gravitational-wave signals encode information about their sources which can be difficult, if not impossible, to otherwise obtain. To extract information from the observed signals requires careful statistical inference. The inferred source parameters can inform our understanding of binary stellar evolution (Stevenson, Ohme & Fairhurst 2015; Abbott et al. 2016d, 2017h; Zevin et al. 2017; Barrett et al. 2018; Belczynski et al. 2018; Bavera et al. 2020), the equation of state of neutron-star matter (Abbott et al. 2018c, 2020b; Most et al. 2018; Essick, Landry & Holz 2020), and the nature of gravity (Yunes & Siemens 2013; Abbott et al. 2016b, 2019h; Yunes, Yagi & Pretorius 2016; Isi et al. 2019). Multimessenger observations of gravitational and electromagnetic radiation (Abbott et al. 2017e) can give an even richer understanding, enabling measurements of cosmological parameters (Abbott et al. 2017d, 2019d; Cantiello et al. 2018; Chen, Fishbach & Holz 2018; Hotokezaka et al. 2019; Dhawan et al. 2020), insights into the structures of gamma-ray bursts (Abbott et al. 2017f; Margutti et al. 2018; Mooley et al. 2018; Fong et al. 2019; Biscoveanu, Thrane & Vitale 2020b), and identifying the origins of heavy elements (Abbott et al. 2017g; Chornock et al. 2017; Tanvir

* E-mail: isobel.romero-shaw@monash.edu

et al. 2017; Kasliwal et al. 2019; Watson et al. 2019). However, electromagnetic emission can fade rapidly, necessitating rapid localization of the gravitational-wave source (Abbott et al. 2018b). To maximize the scientific return of gravitational-wave observations, it is therefore of paramount importance to make use of and continue to develop efficient, reliable, and accurate computational inference.

BILBY is a user-friendly Bayesian inference library that can be used to analyse gravitational-wave signals to infer their source properties (Ashton et al. 2019). BILBY is modular and can be easily adapted to handle a range of inference problems in gravitational-wave astronomy and beyond (e.g. Farah et al. 2019; Goncharov, Zhu & Thrane 2019; Powell & Müller 2019; Sarin, Lasky & Ashton 2020). In the context of gravitational-wave astrophysics and compact binary mergers, it has been used to extract information about short gamma-ray burst properties (Biscoveanu et al. 2020b), neutron star parameters (Coughlin & Dietrich 2019; Hernandez Vivanco et al. 2019a, b; Biscoveanu, Vitale & Haster 2019), the formation history of binary compact objects (Lower et al. 2018; Romero-Shaw, Lasky & Thrane 2019; Ramos-Buades et al. 2020; Romero-Shaw et al. 2020a; Zevin et al. 2020), population properties using hierarchical inference (Abbott et al. 2019i; Talbot et al. 2019; Galaudage, Talbot & Thrane 2019; Kimball et al. 2020), and test general relativity (Keitel 2019; Payne, Talbot & Thrane 2019; Zhao, Lin & Chang 2019; Ashton & Khan 2020; Hübner et al. 2020; Wang & Zhao 2020). This paper concentrates on using BILBY to infer the properties of individual signals from compact binary coalescences – the inspiral, merger, and ringdown of binaries composed of neutron stars and black holes.

We outline the developments included in the BILBY software to accurately and efficiently infer the properties of compact binary coalescence (CBC) signals, and demonstrate their validity both through tests using simulated signals and via comparisons to existing observational results. In Section 2, we describe the applications of Bayesian inference to compact binary coalescence events detected in gravitational waves. In Section 3, we focus on the BILBY package, with particular emphasis on improvements made since the publication of Ashton et al. (2019) in Section 3.1. We outline our code validation tests in Section 3.2, and describe the automation of BILBY – allowing for efficient and immediate analysis of gravitational-wave event candidates – in Section 3.3. In Section 4, we reanalyse the 11 signals from GWTC-1, ensuring that we use both identical data and identical data processing techniques as used to produce the public GWTC-1 results obtained using the Bayesian parameter estimation package LALInference (Veitch et al. 2015). We cross-validate our results for GWTC-1 against these previous results. We defer analysis of detections from the third observing run in anticipation of a future BILBY catalogue. Results of the analyses presented here, in a format matching recent releases of LIGO–Virgo posterior samples, are provided as accompaniments to this paper. Our investigations confirm the effectiveness of BILBY as it begins to be used for LIGO–Virgo parameter estimation (Abbott et al. 2020a, d). Throughout this paper, we use notations for CBC source parameters that are defined in Appendix E.

2 BAYESIAN INFERENCE FOR COMPACT BINARIES

In this section, we outline the fundamental procedures carried out by BILBY and provide a summary of new features implemented since the first BILBY paper (Ashton et al. 2019). For a thorough and up-

to-date description of BILBY, the reader is directed to the BILBY documentation.¹

2.1 Applications of Bayesian inference to compact binary coalescences

The primary objective of gravitational-wave inference for compact binary merger signals is to recover posterior probability densities for the source parameters θ (defined in Appendix E), like the masses and spins of the binary components, given the data and a model hypothesis. The posterior can be computed using Bayes’ theorem (Bayes 1763)

$$p(\theta|d, \mathcal{H}) = \frac{\mathcal{L}(d|\theta, \mathcal{H})\pi(\theta|\mathcal{H})}{\mathcal{Z}(d|\mathcal{H})}, \quad (1)$$

where $\mathcal{L}(d|\theta, \mathcal{H})$ is the likelihood, $\pi(\theta|\mathcal{H})$ is the prior, $\mathcal{Z}(d|\mathcal{H})$ is the evidence, and $\mathcal{H}$ is the model. The prior is chosen to incorporate any *a priori* knowledge about the parameters. The likelihood represents the probability of the detectors measuring data d , assuming a signal (described by the model hypothesis $\mathcal{H}$) with source properties θ . The evidence, or marginalized likelihood

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

serves as a measure of how well the data is modeled by the hypothesis; it acts as a normalization constant in parameter estimation, but is important in model selection.

The standard likelihood function used to analyse gravitational-wave transients is defined in, e.g. Finn (1992) and Romano & Cornish (2017), where both the data and the model are expressed in the frequency domain. This likelihood has stationary Gaussian noise, which is a good approximation in most cases (e.g. Berry et al. 2015; Abbott et al. 2017a, 2019a) unless one of the instruments is affected by a glitch (Pankow et al. 2018; Powell 2018). We assume the noise power spectral density (PSD) is independent of the model parameters and therefore ignore the normalization term, yielding

$$\ln \mathcal{L}(d|\theta) \propto -\sum_k \frac{2|d_k - h_k(\theta)|^2}{TS_k}, \quad (3)$$

where k is the frequency bin index, S is the PSD of the noise, and T is the duration of the analysis segment. The data d and waveform model $h(\theta)$ are the Fourier transforms of their time-domain counterparts. Given the likelihood and the prior, we can calculate the posterior probability distribution for the source parameters.

There are multiple approaches to calculating the posterior probability distribution. For example, RAPIDPE (Pankow et al. 2015) and its iterative spin-off RIFT (Lange, O’Shaughnessy & Rizzo 2018) use highly parallelized grid-based methods to compute the posterior probability distribution, while BAYESTAR (Singer & Price 2016; Singer et al. 2016) rapidly localizes gravitational-wave sources, calculating probabilities on a multiresolution grid of the sky. Bayesian inference schemes using various machine-learning algorithms are also being developed (George & Huerta 2018; Gabbard et al. 2019). However, the majority of Bayesian inference analysis is done by stochastically sampling the posterior probability distribution.

Over many years, Markov-chain Monte Carlo (MCMC; Christensen & Meyer 1998, 2001; Röver, Meyer & Christensen 2006, 2007; van der Sluys et al. 2008b, a) and nested sampling (Veitch & Vecchio 2008, 2010) algorithms for gravitational-wave inference have been developed. This work culminated in the development of LALInference, a Bayesian inference library using custom-built

¹lscsoft.docs.ligo.org/bilby/.

Markov-chain Monte Carlo and nested sampling algorithms (Veitch et al. 2015).² LALINFERENCE has been the workhorse of gravitational-wave inference since the initial LIGO–Virgo era (Aasi et al. 2013), through the first observation (Abbott et al. 2016c) to the production of GWTC-1 (Abbott et al. 2019f). Other stochastic sampling packages used for gravitational-wave inference include PYCBCINFERENCE (Bierwer et al. 2019) and Zackay, Dai & Venumadhav (2018), which uses relative-binning (Cornish 2010; Cornish & Shuman 2020) to reduce the computational cost of the likelihood. In addition to these sampling packages which fit CBC waveform templates to the data, BAYESWAVE (Cornish & Littenberg 2015) uses a trans-dimensional MCMC to fit an *a priori* unknown number of sine-Gaussian wavelets to the data. BAYESWAVE also implements the BAYESLINE algorithm (Littenberg & Cornish 2015) to generate a parametrized fit for the interferometer noise PSD. Power spectral densities produced by BAYESLINE are widely used in gravitational-wave parameter estimation and are used in this work. BILBY has been designed to adapt to the changing needs of the gravitational-wave inference community, emphasizing modularity and ease of accessibility.

While LALINFERENCE implements customized stochastic samplers, BILBY employs external, off-the-shelf samplers, with some adaption. This allows the user to easily switch between samplers with minimal disruption: a useful feature for cross validating results using different samplers. Typically, external samplers need to be tuned and adapted for use in gravitational-wave inference. In some cases, this is a simple case of choosing sensible settings; we provide details of the settings that have been verified for gravitational-wave analysis in Section 4 and Appendix B. However, we also find cases where the off-the-shelf samplers themselves need to be adjusted. Where possible, we propagate those proposed changes to the original sampling packages. Alternatively (e.g. when the change is perhaps gravitational-wave specific), we adjust the sampler from within BILBY.

2.2 Stochastic sampling

Various Monte Carlo sampling schemes have been developed to solve the Bayesian inference problem and estimate the posterior distribution described by equation (1). For low-dimensional problems, a solution might be to estimate the best-fitting parameters by computing the posterior probability for every point on a grid over the parameter space. However, as the number of dimensions increases, this becomes exponentially inefficient.³ The common alternative to solve this problem has been to use stochastic samplers, which fall broadly into two (not mutually exclusive) categories: MCMC (Metropolis et al. 1953; Hastings 1970) and nested sampling (Skilling 2006). In general terms, independent samples are drawn *stochastically* from the posterior, such that the number of samples in the range $(\theta, \theta + \Delta\theta)$ is proportional to $p(\theta|d, \mathcal{H})\Delta\theta$.

²In this work, we focus on Bayesian inference for ground-based gravitational-wave detection. Similar techniques have been developed for studying the gravitational-wave observations of other instruments, such as pulsar timing arrays (Lentati et al. 2014; Vigeland & Vallisneri 2014) and future space-based detectors (Babak et al. 2008, 2010; Marsat, Baker & Dal Canton 2020).

³Quasi-circular binary black hole coalescence waveform models typically have $n_{\text{dim}} = 15$, depending on the number of spin orientations included in the waveform model. Binary neutron star coalescence models include an additional two parameters that describe their tides. We provide definitions of all parameters describing binary compact objects in Appendix E. There are a further ≈ 20 parameters per interferometer that describe uncertainties in detector calibration.

MCMC methods generate posterior samples by noting the positions of particles undergoing a biased random walk through the parameter space, with the probability of moving to a new point in the space given by the transition probability of the Markov chain. Sampling is completed once some user-specified termination condition is reached, usually a threshold for the number of posterior samples that should be accumulated to provide an accurate representation of the posterior.

Nested sampling methods generate posterior samples as a byproduct of calculating the evidence integral $\mathcal{Z}(d|\mathcal{H})$. A set of live points is drawn from the prior distribution, and at each iteration, the live point with the lowest likelihood is replaced by a new nested sample that lies in a part of the parameter space with a higher likelihood. The evidence is approximated by summing the products of the likelihood at the discarded point and the difference in the prior volume between successive iterations. The nested samples are converted to posterior samples by weighting by the posterior probability at that point in the parameter space. The nested sampling algorithm stops once a pre-defined termination condition has been reached. The most commonly used termination condition is when the fraction of the evidence in the remaining prior volume is smaller than a pre-defined amount.

For more details on both MCMC and nested sampling methods, we refer the reader to Hogg & Foreman-Mackey (2018) and Speagle (2020), respectively.

3 THE BILBY PACKAGE

BILBY has a modular structure, allowing users to extend and develop it to suit their needs; examples include online BILBY (Section 3.3.3), BILBY_PIPE (Section 3.3) and parallel BILBY (PBILBY; Section 3.3.2; Smith et al. 2019), amongst others (e.g. Talbot et al. 2019). BILBY comprises three main subpackages. The `core` subpackage contains the basic implementation of likelihoods, priors, sampler interfaces, the result container class, and a host of utilities. The `gw` subpackage builds on `core` and contains gravitational-wave specific implementations of priors and likelihoods. These implementations include a detailed detector and calibration model, an interface to waveform models, and a number of utilities. Finally, the `hyper` subpackage implements hyper-parameter estimation in BILBY, which in the gravitational-wave context is used for population inference.

3.1 Changes within BILBY

Since the original BILBY paper (Ashton et al. 2019), there have been a number of significant changes and added features to the code package. We describe these in the following subsections. We discuss prior constraints in Section 3.1.1, conditional priors in Section 3.1.2, and the implementation of cosmological priors in Section 3.1.3. We detail the custom jump proposals implemented for the CPNEST (Veitch et al. 2017) and PTMCMC (Ellis & van Haasteren 2017) samplers in Section 3.1.5, and the various available prior boundary conditions in Section 3.1.6. Sampling processes can be accelerated using likelihood marginalizations and reduced-order quadratures; we explore how these methods can be applied to BILBY analyses in Sections 3.1.8 and 3.1.9, respectively. In Section 3.1.10, we explain how uncertainties in detector calibration are folded into BILBY parameter estimation. Finally, in Section 3.1.11 we present some of the gravitational-wave transient-specific plots that BILBY can create. In addition to the changes described below, BILBY now also supports the KOMBINE (Farr & Farr 2015), PTMCMC (Ellis & van Haasteren 2017), POLYCHORD (Handley, Hobson & Lasenby 2015a, b), and

ULTRANEST (Buchner 2016, 2019) samplers. A full and up-to-date list of changes can be found in the BILBY changelog.⁴

3.1.1 Constrained priors

Each time the sampler chooses a new point to test from the multidimensional parameter space, it selects this point from within the region specified by the multidimensional prior. It is often advantageous to be able to cut out parts of the prior space by placing restrictions on relationships between parameters. For example, in gravitational-wave inference we frequently wish to specify a prior on the binary component masses, m_1 and m_2 , while enforcing that $m_1 \geq m_2$, which is equivalent to the constraint that the mass ratio $q = m_2/m_1 \leq 1$.

In BILBY, the collection of priors on all parameters is stored as a `PriorDict` object. In order to enforce a constraint, a BILBY user can add a `Constraint` prior object to the `PriorDict`. It is necessary to tell the `PriorDict` how to convert between its sampled parameters and its constrained parameters; this is done by passing a `conversion_function` at instantiation of the `PriorDict`. The BILBY default binary black hole and binary neutron star prior set classes (`BBHPriorDict` and `BNSPriorDict`, respectively) can impose constraints on any of the known binary parameters. This ensures that users can sample in the set of parameters that best suits their problem, while ensuring that the relevant indirectly sampled quantities are constrained. Without applying any prior constraints, all BILBY prior distributions are correctly normalized. When constraints are imposed on the prior distribution, the updated normalization is approximated using a Monte Carlo integral.

3.1.2 Conditional priors

One may choose to make the prior for one parameter conditional on the value of another. This can increase efficiency, particularly if large parts of the prior space would be forbidden by an equivalent constraint prior. A commonly used parametrization of the population distribution of binary black hole masses is

$$p(m_1|m_{\min}, m_{\max}, \alpha) = (1 - \alpha) \frac{m_1^{-\alpha}}{m_{\max}^{1-\alpha} - m_{\min}^{1-\alpha}},$$

$$p(q|m_1, m_{\min}, \beta) = (1 + \beta) \frac{m_1^{1+\beta} q^\beta}{m_1^{1+\beta} - m_{\min}^{1+\beta}}, \quad (4)$$

where $m_{\min}$ and $m_{\max}$ are the maximum and minimum allowed masses for the primary component, and α and β are power-law indices (Fishbach & Holz 2017; Abbott et al. 2019i). If we wish to use a similar prior to analyse individual binary black hole coalescences, we require a prior for mass ratio which is conditioned on the primary mass. We provide a `ConditionalPriorDict` and conditional versions of all implemented priors within BILBY to facilitate analyses of this kind. Further, BILBY is able to handle nested and multiple dependencies, and automatically resolves the order in which conditional priors need to be called. The conditional relationship between different priors can have any functional form specified by the user.

3.1.3 Cosmological priors

Most previous parameter estimation analyses of CBCs have assumed a prior on luminosity distance d_L which is $\pi(d_L) \propto d_L^2$ (e.g. Abbott

Figure 1. Comparison of distance priors out to redshift $z = 0.10$ (top panel) and $z = 1.02$ (bottom panel), respectively, corresponding to $d_L = 500$ and $d_L = 7000$ Mpc, according to Ade et al. (2016) cosmology. The upper and lower panels show the range of the luminosity distance priors for the default 128 s and high-mass prior sets, respectively. We display priors that are uniform in luminosity volume, comoving volume, and the (comoving) source frame. The probability density of each curve is normalized with respect to the upper limit cut-off displayed in that panel.

et al. 2016c, 2019f). A $\pi(d_L) \propto d_L^2$ prior would distribute mergers uniformly throughout a Euclidean universe. This is an adequate approximation at small redshifts, as illustrated in Fig. 1; however, beyond a redshift of ~ 1 , the difference between a prior which is uniform in the comoving (source) frame volume and uniform in luminosity volume is large. We therefore implement a range of cosmologically informed prior classes.

The `Cosmological` base class allows the user to specify a prior in either luminosity distance, or comoving distance, or redshift using any cosmology supported in `Astropy` (Robitaille et al. 2013; Price-Whelan et al. 2018).⁵ Additionally, users can specify the prior in terms of redshift and then convert to an equivalent prior on luminosity distance if desired. We implement two new source distance priors: a `UniformComovingVolume` prior, defined as

$$\pi(z) \propto \frac{dV_c}{dz}, \quad (5)$$

where V_c is the comoving volume, and a `UniformSourceFrame` prior, defined as

$$\pi(z) \propto \frac{1}{1+z} \frac{dV_c}{dz}. \quad (6)$$

The additional factor of $(1+z)^{-1}$ accounts for time dilation.

Additional `Cosmological` prior classes of the form

$$\pi(z) \propto \frac{dV_c}{dz} f(z) \quad (7)$$

can be defined by providing $f(z)$.

⁴git.ligo.org/lscsoft/bilby/blob/master/CHANGELOG.md.

⁵By default, BILBY uses the Ade et al. (2016) cosmology.

3.1.4 Joint priors

In cases where one requires more complex priors that depend on multiple parameters we implemented the `JointPrior` class in which the user can define a distribution that describes the prior on multiple parameters. This is implemented in BILBY in the `MultivariateGaussian` prior that lets the user define multimodal and multivariate Gaussian priors. It is also used in the `HEALPIXMAP` prior in which a user can implement a prior on the sky position and optionally distance according to a given HEALPIX (Górski et al. 1999, 2005) map.

3.1.5 Custom jump proposals

Users of BILBY can define custom jump proposals through its interface to the CPNEST and PTMCMC samplers. Jump proposals describe how the sampler finds new points in the parameter space. CPNEST has a defined cycle of proposals that can be changed by the user. These proposals can be useful when there are known degeneracies in the parameter space, e.g. phase ϕ and polarization angle ψ under a shift by $\pi/2$ in either parameter (Veitch et al. 2015). Sampling in right ascension α and declination δ can also be improved using custom jump proposals; degeneracy typically leads to a ring-shaped 2D posterior in these parameters for signals detected by two detectors (Singer et al. 2014; Berry et al. 2015). We provide proposals for the above two cases in the BILBY implementation of CPNEST, while additional proposals can be defined by the user to suit their needs.

3.1.6 Boundary conditions

For many parameters, such as the mass ratio q and spin magnitudes a_1, a_2 , posterior distributions have significant support close to the prior boundaries. This is expected behaviour and a direct result of the choice of prior (e.g. the choice to fix $m_1 \geq m_2$ ensures $q \leq 1$). In BILBY, `Prior` objects have boundaries that can be specified by the user as `None`, `reflective`, or `periodic`. For samplers that support these settings, these options specify the behaviour of the sampler when it proposes a point that is outside of the prior volume. For a `None` boundary, such a point is rejected. Priors that have `reflective` boundaries are reflected about the boundary (a proposed mass ratio of $1 + \epsilon$ is reflected to $1 - \epsilon$) while `periodic` boundaries wrap around (a proposed phase of $\pi + \epsilon$ is wrapped to ϵ).

The DYNesty sampler (Speagle 2020) supports all available parameters boundary settings. The PYMULTINEST sampler (Feroz & Hobson 2008; Feroz, Hobson & Bridges 2009; Buchner et al. 2014; Feroz et al. 2019) can implement `periodic` boundary conditions, but not `reflective`, which are treated as `None`. All other samplers implemented in BILBY treat all prior boundaries as `None`.

While `reflective` boundaries are implemented, their usage is not recommended due to concerns that they break detailed balance (e.g. Suwa & Todo 2010). When using the DYNesty sampler, we recommend using `periodic` boundaries for relevant parameters (e.g. the right ascension and phase). These recommendations are mirrored in our choices of default priors, discussed in Section 4.1.

3.1.7 Alternative sky and time parameterisations

The most common way to describe the location of the source on the sky and its time of arrival is with the equatorial coordinates right ascension α and declination δ , and the coalescence time at the centre of the Earth t_c . However, particularly when the signal is only

observed in two detectors, the likelihood is determined primarily by the time delay between the arrival of the signal at each detector. The posterior distribution on these parameters often assumes a broken ring shape misaligned with the equatorial coordinate system (Singer et al. 2014; Berry et al. 2015), making sampling difficult. A more natural parametrization of the problem is given by sampling in the time of arrival at one of the detectors [ideally the one with the largest signal-to-noise ratio (SNR)], and rotating the sky coordinates such that the ring structure is uncorrelated in the sampling parameters.

We allow the user to specify a `reference_frame` and `time_reference`. The argument `reference_frame` can either be an `InterferometerList`, a string with the names of two known detectors, e.g. `H1L1`, or `sky` to sample in α and δ . Cases where sampling in α and δ is preferred include when the astrophysical location of the source is exactly known, e.g. by using the location of the host galaxy of a binary neutron star merger, the user can sample in α and δ by specifying `reference_frame=sky`. In this parametrization the zenith angle κ is related to the time delay of the merger between the two detectors and is therefore well measured. The azimuthal angle ϵ is only weakly constrained for a two-detector network. The argument `time_reference` can be the name of any known interferometer, e.g. `H1`, or `geocent` to sample in the time at the geocentre.

The detector-based sampling frame is defined in terms of the zenith κ and azimuthal ϵ angles relative to the vector connecting the vertices of the two interferometers specified δr . We perform the transformation from (κ, ϵ) to (δ, α) by constructing the rotation matrix R which maps $\hat{z}$ to the unit vector $\delta \hat{r}$. The rotation matrix R can be described by three Euler angles (α, β, γ)

$$R = R_3(\gamma)R_2(\beta)R_3(\alpha),$$

$$\tan \alpha = \frac{-\delta r_y \delta r_z}{\delta r_x}, \quad \cos \beta = \delta r_y, \quad \tan \gamma = \frac{\delta r_y}{\delta r_x}. \quad (8)$$

Here, $\delta r_{\{x,y,z\}}$ are the Cartesian components of δr and $R_{2,3}$ are rotation matrices about the y - and z -axes respectively.

3.1.8 Analytical likelihood marginalizations

The likelihood in equation (3) can be costly to evaluate for some signal models, and the size of the coalescence-time posterior relative to its much wider prior can make sampling the entire space difficult. Therefore, we reduce the dimensionality of the CBC problem by analytically marginalizing over certain parameters, speeding up computation and improving the sampler convergence. The parameters we commonly marginalize over are the coalescence time, binary orbital phase, and luminosity distance. In the frequency domain, a waveform of total duration T can be written in terms of a reference time t_0 , phase ϕ_0 , and luminosity distance d_0 as

$$h_k(\lambda, t, \phi, d_L) = h(\lambda, t_0, \phi_0 = 0, d_0) \times \exp \left[-2\pi i k \frac{(t - t_0)}{T} \right] \exp(2i\phi) \frac{d_0}{d_L}, \quad (9)$$

where k indicates the frequency bin and λ represents the set of the other binary parameters, including the masses and spins, whose contributions to the waveform cannot be separated and thus cannot be analytically marginalized. The phase dependence can only be factored out for waveforms that include just the dominant $\ell = 2, m = 2$ mode; however, this factorization has been shown to be a reasonable approximation in some cases when precession is not measurable (Abbott et al. 2017b). The marginalized likelihood is obtained by integrating the likelihood in equation (3) over phase,

distance, and coalescence time after using the factorisation in equation (9). The phase integral simplifies to a modified Bessel function of the first kind, evaluated at the magnitude of the complex inner product of the waveform and the data (Veitch & Del Pozzo 2013; Veitch et al. 2015).

The distance marginalization is performed numerically, using a Riemann sum in matched filter and optimal SNR over the range $\rho \in [10^{-5}, 10^{10}]$, spaced uniformly in log-space (Singer & Price 2016; Singer et al. 2016; Thrane & Talbot 2019). To improve efficiency at run-time, we build a lookup table which is interpolated and then evaluated. The lookup table is computed before the sampling phase begins, and can be cached and reloaded from previous analyses that used the same distance prior.

The marginalization over time involves performing a quadrature integral over an evenly spaced array of times separated by the sampling frequency. This marginalization is enabled by the fact that the inner product of the time-domain waveform and data can be rewritten as a fast Fourier transform (Farr 2014). The sky location inferred when sampling in the sky frame and using the time-marginalized likelihood is not generally correct and we do not recommend combining these two features.

If the signal is loud and the sampling frequency is too low, the reconstructed coalescence-time posterior appears discrete, since each of the generated parameters lies on one of the nodes of the array. One solution to this is to increase the resolution of the array times by increasing the sampling frequency. However, this increases the computational cost of the marginalized likelihood evaluation. Additionally, gravitational-wave detector data are natively sampled at 16 kHz (Abbott et al. 2019e), so increasing the time resolution beyond this level would require a different technique, e.g. zero-padding. In order to avoid increasing the sampling frequency, we maintain a continuous coalescence-time posterior by introducing a `time_jitter` δt . This parameter varies the position of the time array over which the numerical integral is performed. We apply a uniform prior with bounds such that

$$-\frac{T}{2} \leq \delta t < \frac{T}{2}, \quad (10)$$

thus reducing the prior space to be searched.

When using the analytically marginalized likelihood, the sampler does not produce posterior samples for the marginalized parameters. However, BILBY is able to generate samples for these parameters in post-processing. Using BILBY, we recalculate the likelihood by recomputing the optimal matched filter SNR and the inner product of the waveform and data. We then obtain a posterior array for the marginalized parameter in question, evaluated at discrete points in the parameter's prior space. We generate posterior samples by sampling from this interpolated posterior array. By drawing a single sample for each of the marginalized parameters for each posterior sample we maintain the degeneracies between, e.g. distance and binary orbital inclination. For detailed derivations of the analytically marginalized likelihood and the posterior sample reconstruction process (see Thrane & Talbot 2019).

3.1.9 Reduced-order quadrature

In order to reduce the number of frequencies at which the likelihood in equation (3) must be evaluated, we implement the reduced-order quadrature (ROQ) likelihood (Smith et al. 2016). This method works by identifying a reduced basis that can describe the signal model well over a certain range of the parameter space. Application of reduced-order methods has been crucial for expediting inference for long

duration signals, such as the binary neutron star merger GW170817 (Abbott et al. 2019f). Evaluating the ROQ likelihood requires access to the appropriate basis. A set of bases for the most commonly used waveform, IMRPHENOMPv2, are publicly available online.⁶

The ROQGravitationalWaveTransient likelihood class in BILBY is able to analyse arbitrary reduced-order bases. This likelihood can also be marginalized over phase and/or distance. A time-marginalized ROQ likelihood has not yet been implemented.

3.1.10 Calibration

The imperfect nature of the detector calibration introduces a systematic error in the measured astrophysical strain (Abbott et al. 2016d). Following Farr, Farr & Littenberg (2014), we split this error into frequency-dependent amplitude and phase offsets, $\delta A(f)$ and $\delta\phi(f)$, respectively. The observed strain can then be related to the true strain as

$$h_{\text{obs}}(f) = h(f) [1 + \delta A(f)] \exp[i\delta\phi(f)]. \quad (11)$$

Since the calibration error is small, we perform a small angle expansion in the phase correction

$$\exp[i\delta\phi(f)] = \frac{2 + i\delta\phi(f)}{2 - i\delta\phi(f)} + \mathcal{O}(\delta\phi^3). \quad (12)$$

Substituting this, we obtain

$$h_{\text{obs}}(f) = h(f) [1 + \delta A(f)] \frac{2 + i\delta\phi(f)}{2 - i\delta\phi(f)}. \quad (13)$$

The amplitude and phase uncertainty are modelled as cubic splines in BILBY

$$\delta A(f) = s(f; \{f_j, \delta A_j\}), \quad (14)$$

$$\delta\phi(f) = s(f; \{f_j, \delta\phi_j\}), \quad (15)$$

where the spline nodes f_j are fixed and distributed uniformly in log-space between the minimum and maximum frequencies included in the likelihood, and the values of the splines at the nodes, δA_j and $\delta\phi_j$, are sampled parameters (Vitale et al. 2012).

The priors on the spline values are taken to be normal distributions, with means and widths that can either be constant or loaded from a frequency-dependent calibration envelope file (Cahillane et al. 2017; Viets et al. 2018). The calibration factor defined in equations (14) and (15) are applied to the waveform calculated for each prior sample before the likelihood is computed. Fig. 2 shows an example plot of the calibration spline posterior for both the amplitude and phase uncertainties.

3.1.11 Gravitational-wave transient-specific plots

BILBY users can produce sets of posterior plots specific to gravitational-wave transient analysis. We use the `ligo.skymap` (Singer & Price 2016; Singer et al. 2016) package to produce sky maps in both the `fits` format commonly used for electromagnetic observation and standard image formats. We are also able to produce plots showing our inferred posterior on the detector calibration and waveform models, in addition to the parameters describing these models. We present examples of these plots for GW150914 in Figs 2 and 3, respectively. In such plots, we show the mean reconstructed model and symmetric 90 per cent credible intervals.

⁶git.ligo.org/lscsoft/ROQ_data.

Figure 2. Calibration posteriors for the amplitude (top) and the phase uncertainty (bottom) for both LIGO Hanford (orange) and Livingston (blue) detectors for GW150914. The solid curves show the mean, while the shaded region represents the 90 per cent confidence intervals. The vertical lines show the locations of the spline points.

Figure 3. Reconstructed waveform for GW150914 for LIGO Hanford. The top panel shows the amplitude spectral density of the signal (blue), data (light orange), and estimated noise amplitude spectral density (dark orange). The bottom panel shows the time domain data (light orange) and waveform estimate (blue). The dark blue curves show the mean recovered waveform and the light blue shaded region the 90 per cent confidence interval.

3.2 Validation of BILBY

A common consistency test of the performance of sampling algorithms is to check that the correct proportion of true parameter values are found within a given probability interval for simulated systems (Cook, Gelman & Rubin 2006; Talts et al. 2018) – i.e. that 10 per cent of events are found within the 0.1 probability credible interval, 50 per cent are found within the 0.5 probability credible interval, etc. We generate a set of CBC signals with true parameter

Figure 4. Results of 100 injections drawn from the four-second prior defined in Section 4.1. The grey regions cover the cumulative 1σ , 2σ , and 3σ confidence intervals in order of decreasing opacity. Each coloured line tracks the cumulative fraction of events within this confidence interval for a different parameter. The combined p -value for all parameters, over all tests, is 0.7206, consistent with the individual p -values being drawn from a uniform distribution. Individual parameter p -values are displayed in parentheses in the plot legend. The marginalized parameters – geocentre time t_c , luminosity distance d_L , and phase ϕ – are reconstructed in post-processing. Other parameters provided in the plot legend are defined in Appendix E.

values drawn from our prior probability distributions and inject these into simulated noise. Parameter estimation is then performed on each signal to determine the credible level at which the true value of each parameter is found. This test is traditionally used in validating gravitational-wave inference codes (Sidery et al. 2014; Berry et al. 2015; Pankow et al. 2015; Veitch et al. 2015; Singer & Price 2016; Del Pozzo et al. 2018; Biwer et al. 2019).

To test BILBY’s parameter estimation, we simulate 100 synthetic CBC signals for a two-detector Hanford–Livingston network and add the signals to Gaussian noise coloured to the anticipated Advanced LIGO design sensitivity (Abbott et al. 2018b). The parameters of the simulated events are drawn from the default 4 s prior set, detailed in Section 4.1.

Parameter estimation is performed using the DYNesty sampler with the distance, time, and phase-marginalized likelihood. Analysis of the performance of other samplers is left to future work. Results of the test are shown in Fig. 4, where the fraction of events for which the true parameter is found at a particular confidence level is plotted against that particular confidence interval.⁷ We also show the individual parameter p -values representing the probability that the fraction of events in a particular confidence interval is drawn from a uniform distribution, as expected for a Gaussian likelihood, and the combined p -value quantifying the probability that the individual p -values are drawn from a uniform distribution. The combined p -value obtained with the latest version of BILBY is 0.7206 and the minimum is 0.183 for ϕ , which is entirely consistent with chance for the set of 15 parameters, indicating that the posterior probability distributions

⁷These plots are referred to as P–P plots, where P could stand for probability, percent or proportion. Instructions for generating P–P plots are provided in the BILBY documentation at git.ligo.org/lscsoft/bilby_pipe/wiki/pp/howto.

produced by BILBY are well-calibrated. The grey regions show the 1σ , 2σ , and 3σ confidence intervals so we expect the lines to deviate from this region approximately 0.3 per cent of the time, which is consistent with what we see.

In addition to the procedure described above, we verify the suitability of the sampler settings for the problem of sampling the CBC parameter space using a series of review tests. These are described in detail in Appendix A. The settings used for each of the tests described here are provided in Appendix B. In addition to these review tests, BILBY has an extensive set of unit tests, which scrutinize the behaviour of the software in high detail every time a change is made to the code; these unit tests can be found within the BILBY package.⁸

3.3 Automation of BILBY for gravitational-wave inference

With the improvement in sensitivity and expansion of the gravitational-wave observatory network comes an increasing rate of detections. Streamlining the deployment of BILBY analysis is therefore vital. We introduce BILBY_PIPE, a PYTHON package providing a set of command-line tools designed to allow performance of parameter estimation on gravitational-wave data with all settings either passed in a configuration file or via the command line.⁹ This tool was used to perform the analyses of the GWTC-1 catalogue events presented in Section 4, and is integral to the automatic online parameter estimation that is triggered by potential gravitational-wave events.

The BILBY_PIPE workflow consists of two key stages: data generation and data analysis. These steps are outlined in Section 3.3.1. The pipelines provided by BILBY_PIPE can be utilized to distribute analysis of a single event over multiple CPUs using PBILBY (Smith et al. 2019), which is described in Section 3.3.2. The workflow for the automated running of BILBY on gravitational-wave candidates is detailed in Section 3.3.3.

3.3.1 Data generation and analysis

Gravitational-wave detectors record and store time-domain strain data and information about the behaviour internal to the detectors, as well as data from a suite of environmental sensors. To obtain gravitational-wave strain data, we recommend using the GWPY library (Macleod et al. 2018). GWpy can retrieve both public data from the Gravitational Wave Open Science Center (Abbott et al. 2019e), and proprietary data using the Network Data Server protocol (NDS2) to acquire data from LIGO servers. Given a GPS trigger time and a required data duration, BILBY_PIPE uses GWPY to extract an analysis segment of strain data around the trigger, as well as a segment of strain data used to estimate the noise PSD. The default duration for the analysis segment is $T = 4$ s, which is considered adequate for sources with detector-frame chirp masses $\mathcal{M} \gtrsim 15 M_{\odot}$. Sources with lower $\mathcal{M}$ have longer signals, so longer analysis segments should be used. A portion of data following the trigger time is required to encompass the remaining merger and post-coalescence ringdown signal; this is 2 s by default.

A BILBY_PIPE user can provide pre-generated PSDs, and a range of design-sensitivity noise spectra for current and future detectors is available as part of the BILBY package. For the analyses we

present in Section 4, we use event-specific PSDs produced using BAYESWAVE (Cornish & Littenberg 2015). When a PSD is not provided, BILBY_PIPE uses the median-average power spectrum method described by Allen et al. (2012), and implemented in GWpy, to calculate the PSD; this method has the advantage of downweighting outliers in the off-source data (Allen et al. 2012; Veitch et al. 2015). In order to avoid including any signal in the PSD calculation, BILBY_PIPE uses a stretch of data preceding the analysis segment. Following Veitch et al. (2015) and Chatziioannou et al. (2019), we use data stretches of length $\min(32T, 1024 \text{ s})$ by default, although both of these values can be altered by the user. The upper limit of 1024 s is required because the PSD of gravitational-wave detectors is non-stationary over long time-periods (Chatziioannou et al. 2019). To further mitigate this issue, the data are divided into segments of length T , with each segment overlapping 50 per cent of the previous segment; this allows a shorter total stretch of data to be used to calculate the PSD. Following Allen et al. (2012), segments are Tukey windowed with a 0.4 s roll-off to suppress spectral leakage (Abbott et al. 2019a), before computing their one-sided power spectra.

The priors for the analysis can be specified by the user, either by providing a path to a file containing the priors in BILBY syntax, or by giving the name of one of the default BILBY_PIPE priors described in Section 4.1. By default, the BILBY GravitationalWaveTransient likelihood is used with the waveform template generated by LALSIMULATION (LIGO Scientific Collaboration 2018). However, users can specify their own source models and modified likelihoods in the configuration file. After saving the necessary data, BILBY_PIPE launches parameter estimation on the analysis segment in accordance with the procedure outlined in Section 2.1.

3.3.2 Parallel BILBY

Parallel BILBY (Smith et al. 2019) is a parallel implementation of BILBY which uses Message Passing Interface (MPI; Farah et al. 2019) to distribute the DYNESTY nested sampling package over a pool of CPUs. Nested sampling requires drawing successive samples satisfying a likelihood constraint from the prior. Faithfully drawing samples from this constrained prior requires many likelihood evaluations. We use a CPU pool to draw prior samples in parallel at each iteration of the algorithm to reduce the wall-time needed to complete an analysis.

Qualitatively, PBILBY works by using a pool of n_{cores} CPUs to draw $n_{\text{cores}} - 1$ samples from the prior in parallel at each iteration of the sampling algorithm. The $n_{\text{cores}} - 1$ proposed samples are ranked by likelihood and the lowest likelihood live point is replaced. The prior volume is then updated on all n_{cores} processes and the sampling step is repeated until the algorithm is converged. The speedup S of the parallel implementation is a function of the number of live points n_{live} and the number of parallel processes (Smith et al. 2019)

$$S = n_{\text{live}} \ln \left(1 + \frac{n_{\text{cores}}}{n_{\text{live}}} \right). \quad (16)$$

Currently, PBILBY only supports the DYNESTY and PTMCEE sampling packages. All of the functionality of BILBY, as described in Section 3.1, is supported by PBILBY.

PBILBY is highly scalable, and is thus well suited to accelerating applications in which the gravitational-wave signal or noise models are computationally expensive to evaluate, e.g. time-domain signal models such as spin-precessing effective-one-body models with higher order modes (Bohé et al. 2017; Ossokine et al. 2020), numerical-relativity surrogate models (Blackman et al. 2017) and models including tidal effects (Nagar et al. 2018; Lackey et al. 2019).

⁸git.ligo.org/lscsoft/bilby/tree/master/test.

⁹The source-code is available on the git repository git.ligo.org/lscsoft/bilby_pipe. Specifics about the installation, functionality, and user examples are also provided lscsoft.docs.ligo.org/bilby_pipe.

Other well-suited applications include those where sampling convergence can be slow due to high dimensionality of the parameter space, e.g. when calibration (Farr et al. 2014) or beyond-general-relativity parameters are used (Abbott et al. 2016b, 2019h), or when a large number of live points is required to effectively estimate the evidence.

In order to facilitate efficient inter-CPU communication with MPI, PBILBY is a stand-alone package, though it still uses the underlying BILBY modules.

In addition to the hugely parallel PBILBY, many of the implemented sampling packages support parallelization through a user specified pool of processes. For these samplers BILBY natively supports local parallelization using the PYTHON MULTIPROCESSING package. When available, the number of parallel computational threads to use is specified using the `nthreads` argument.

3.3.3 Online BILBY

The gravitational-wave candidate event data base GraceDB¹⁰ provides a centralized location for collecting and distributing gravitational-wave triggers uploaded in real time from search pipelines. Once uploaded, each trigger is assigned a unique identifier, and LIGO–Virgo users are notified via an LVALERT (LIGO–Virgo Alert Network). GWCELERY (Singer et al. 2020), a PYTHON-based package designed to facilitate interactions with GraceDB, responds to an alert by first creating a Superevent, which groups triggers from multiple search pipelines and then chooses a preferred event based on the SNR of the triggers. If the preferred candidate has a false-alarm-rate (FAR) below a given threshold, GWCELERY automatically launches multiple parameter estimation jobs. For the case of BILBY, this involves making a call to the `bilby_pipe_gracedb` executable.

The `bilby_pipe_gracedb` executable takes the GraceDB event ID as input and generates a configuration file based on the trigger time of the candidate. A prior file is selected from the set of default priors using the chirp mass of the gravitational-wave signal template that triggered the LVALERT. Further details about the default priors can be found in Section 4.1. These files are then passed to the `bilby_pipe` executable, which runs parameter estimation on the event. PESUMMARY (Hoy & Raymond 2020), a PYTHON-based package designed to post-process inference package output in a number of formats, then generates updated source classification probabilities and webpages displaying diagnostic plots. Once this step is complete, GWCELERY uploads the posterior samples, post-processing pages and updated source classification probabilities to GraceDB. Fig. 5 illustrates the process of automated parameter estimation from the trigger of a gravitational-wave event to the upload of BILBY parameter estimation results to GraceDB.

3.3.4 Run times

The overall run time of a BILBY parameter estimation job depends on the specific input data and can vary considerably based on the chosen sampler settings and SNR. The overall wall time can be reduced by allowing for marginalization over certain parameters, as described in Section 3.1.8, or by using the parallelization methods described in Section 3.3.2. For a GW150914-like binary black hole merger, the expected run time for a time, distance, and phase marginalized BILBY analysis using the default waveform model IMRPHENOMPV2 (Schmidt, Hannam & Husa 2012) is $\mathcal{O}(10)$ h. The waveform models

Figure 5. Workflow for online BILBY parameter estimation.

needed to analyse binary neutron star merger events are much longer than those required for binary black holes, and therefore are more computationally expensive. Hence, for a GW170817-like binary neutron star merger event, we use PBILBY to distribute the analysis over a pool of CPUs, as described in Section 3.3.2; the expected run time in this case is $\mathcal{O}(10)$ h.

4 GRAVITATIONAL-WAVE TRANSIENT CATALOGUE

This section contains our run settings for performing parameter estimation on GWTC-1 events using BILBY, in addition to the results we obtain from this analysis. We describe our default priors and sampler settings in Sections 4.1–4.4. Further details about these settings are given in Appendix B. We provide our results in Section 4.6, where we assess their statistical similarity to those published in GWTC-1 (Abbott et al. 2019f).¹¹ All BILBY_PIPE configuration files, posterior samples and BILBY results files are made available online (Romero-Shaw et al. 2020b).

¹¹The LALINFERENCE posterior samples that we show in this section are taken from the Parameter Estimation Sample Release for GWTC-1 (Abbott et al. 2018a). The posterior samples from LALINFERENCE are obtained using a mixture of the nested sampling algorithm of LALInferenceNest and the Markov chain Monte Carlo algorithm of LALInferenceMCMC (Veitch et al. 2015).

¹⁰gracedb.ligo.org.

Table 1. Summary statistics for each event in GWTC-1, as recovered by BILBY. We quote median values along with the symmetric 90 per cent credible interval range around the median. For mass ratio q , we quote the 90 per cent lower limit (10 per cent quantile), with all events being consistent with equal mass ($q = 1$). We use a fixed-sky prior on source location for GW170817, the binary neutron star merger, fixing the source at the right ascension and declination of its electromagnetic counterpart (Abbott et al. 2017e). The 90 per cent credible areas for sky location are computed using 3000 samples from each posterior. The final column lists the maximum JS divergence statistic (a measure of the similarity between two distributions) between the BILBY GWTC1 samples, and the LALINFERENCE GWTC-1 posterior samples across the model parameters. We consider JS divergence values greater than 0.002 bit to be statistically significant.

Event	Prior	$\mathcal{M}/M_\odot$	$\mathcal{M}^{\text{source}}/M_\odot$	q lower limit	d_L/Mpc	χ_{eff}	$\Delta\Omega/\text{deg}^2$	Max-JS/bit
GW150914	4 s	31^{+1}_{-1}	28^{+2}_{-1}	0.72	420^{+160}_{-165}	$-0.0^{+0.1}_{-0.1}$	169	$\text{JS}_{\theta_{JN}} = 0.0019$
GW151012	4 s	18^{+2}_{-1}	15^{+2}_{-1}	0.41	1015^{+498}_{-472}	$0.0^{+0.2}_{-0.2}$	1457	$\text{JS}_{\mathcal{M}} = 0.0014$
GW151226	8 s	$9.7^{+0.1}_{-0.1}$	$8.9^{+0.3}_{-0.3}$	0.38	428^{+196}_{-189}	$0.2^{+0.1}_{-0.1}$	1022	$\text{JS}_q = 0.0017$
GW170104	4 s	26^{+2}_{-2}	22^{+2}_{-2}	0.48	935^{+441}_{-411}	$-0.0^{+0.2}_{-0.2}$	900	$\text{JS}_{\mathcal{M}} = 0.0007$
GW170608	16 s	$8.5^{+0.0}_{-0.0}$	$7.9^{+0.2}_{-0.2}$	0.49	317^{+122}_{-115}	$0.0^{+0.1}_{-0.0}$	1462	$\text{JS}_q = 0.0011$
GW170729	High mass	51^{+8}_{-9}	35^{+6}_{-5}	0.43	2548^{+1369}_{-1235}	$0.3^{+0.2}_{-0.3}$	1050	$\text{JS}_\alpha = 0.0026$
GW170809	4 s	30^{+2}_{-2}	25^{+2}_{-2}	0.51	995^{+311}_{-411}	$0.1^{+0.2}_{-0.2}$	300	$\text{JS}_{\mathcal{M}} = 0.0010$
GW170814	4 s	27^{+1}_{-1}	24^{+1}_{-1}	0.69	572^{+154}_{-212}	$0.1^{+0.1}_{-0.1}$	77	$\text{JS}_{\theta_1} = 0.0009$
GW170817	Custom	$1.1975^{+0.0001}_{-0.0001}$	$1.187^{+0.004}_{-0.002}$	0.74	40^{+8}_{-16}	$0.00^{+0.02}_{-0.01}$	N/A	$\text{JS}_{\bar{\Lambda}} = 0.0019$
GW170818	4 s	32^{+2}_{-2}	27^{+2}_{-2}	0.58	1017^{+407}_{-348}	$-0.1^{+0.2}_{-0.2}$	29	$\text{JS}_\alpha = 0.0064$
GW170823	High mass	39^{+5}_{-4}	29^{+4}_{-3}	0.54	1771^{+857}_{-831}	$0.0^{+0.2}_{-0.2}$	1570	$\text{JS}_{\theta_N} = 0.0009$

Figure 6. Difference between the right ascension (α) samples recovered by BILBY and LALINFERENCE for all BBH events. This is the worst recovered parameter according to the JS divergence. Labels show the mean JS divergence between α samples, evaluated by random re-sampling over 100 iterations.

4.1 Default priors

The default prior distributions contained in BILBY_PIPE are predominantly tailored to specific signal durations, with the exception of a high-mass prior tailored to particularly heavy sources with detector-frame chirp mass $\mathcal{M}$ up to $175 M_\odot$. For each event in GWTC-1, we choose the default prior that best covers the prior volume studied using LALINFERENCE for the original samples release. This means that two events (GW150914 and GW151012) are analysed using priors suited to signals of duration $T = 4$ s, even though we match the data duration to that used in the original LALINFERENCE analysis ($T = 8$ s). The prior on $\mathcal{M}$ is uniform in the detector frame, while the prior on d_L is uniform in comoving volume and source frame time, as implemented in the `UniformSourceFrame` prior class described in Section 3.1.3. The $\mathcal{M}$, d_L , and spin magnitude prior limits vary between prior sets, while the other source parameters are assigned priors that are consistent between sets. The shapes and limits of all

Figure 7. Difference between the luminosity distance (d_L) samples recovered by BILBY and LALINFERENCE for all events. Labels show the mean JS divergence between d_L samples, evaluated by random re-sampling over 100 iterations.

priors are defined in Appendix B2. The prior files can be found in the BILBY_PIPE git repository.⁹

4.2 Likelihood

Our likelihood is marginalized over reference phase and source luminosity distance, as described in Section 3.1.8. For binary black hole merger analyses, we use the waveform model IMRPHENOMPv2 (Schmidt et al. 2012; Hannam et al. 2014; Bohé et al. 2016; Khan et al. 2016) as our signal template. For the binary neutron star GW170817, we use the IMRPhenomPv2_NRTidalv2 waveform model with tidal effects (Dietrich et al. 2019).

4.3 Sampling

We use DYNesty (Speagle 2020) as our sampler; see Appendix B1 for the detailed sampler settings. We use the static version of DYNesty,

Figure 8. Comparison of the posterior distributions between the LALINFERENCE (grey) and BILBY (coloured) packages over the source primary mass m_1^{source} and source secondary mass m_2^{source} parameter space. Each contour shows the 90 per cent credible area, with the LALINFERENCE posterior samples reweighted to the BILBY priors.

Figure 9. Posterior probability distributions for source-frame chirp mass $\mathcal{M}^{\text{source}}$ and luminosity distance d_L for GW150914. We display posteriors obtained using BILBY in orange, and LALINFERENCE posteriors in blue. We reweight the LALINFERENCE posteriors to the BILBY default priors using the procedure outlined in Appendix C. The 1D JS divergence on chirp mass $\mathcal{M}$ and luminosity distance d_L for this event are $\text{JS}_{\mathcal{M}} = 0.0017$ bit and $\text{JS}_{d_L} = 0.0015$ bit.

as is default for BILBY_PIPE. For each event, we run five analyses in parallel, merging the resultant posterior samples in post-processing. When combining results, care must be taken to weight each set of samples appropriately by its relative evidence. The weight applied to the i th component of N sets of posterior samples is given by

$$w_i = \frac{\mathcal{Z}_i}{\sum_{j=1}^N \mathcal{Z}_j}, \quad (17)$$

where $\mathcal{Z}_i$ is the evidence of the i th set of samples.

4.4 Data used

We use detector noise PSDs and calibration envelopes data from the data releases accompanying GWTC-1 (Abbott et al. 2019b, c, f). The data for each event are obtained through BILBY_PIPE using methods from the GWpy (Macleod et al. 2018) package as outlined in Section 3.3.1. Appendix B contains details of the trigger times and data segment durations specified for each event, which we choose to match those used in the original LALINFERENCE analysis.

4.5 Analysis of binary neutron star merger GW170817

The first observation of a binary neutron star coalescence, GW170817, by LIGO–Virgo (Abbott et al. 2017c) presented a new challenge for gravitational-wave transient inference. The longer signal durations increase the typical computing requirements, and for systems containing a neutron star, tidal effects become important in the waveform models. The original discovery (Abbott et al. 2017c) and subsequent follow-up studies (Abbott et al. 2019g) analysed the data with a variety of waveform models and under differing assumptions.

We employ PBILBY for this analysis, with BILBY_PIPE default sampler settings. We use priors chosen to match those of the LVC analysis (Abbott et al. 2019g), but sample in chirp mass and mass ratio rather than component masses. Our likelihood is computed using the tidal waveform model IMRPhenomPv2_NRTidalv2 (Dietrich et al. 2019). This PBILBY analysis took approximately 11 h on 560 cores.

4.6 Results

We make posterior samples and BILBY_PIPE configuration settings files available online (Romero-Shaw et al. 2020b, c). To directly compare BILBY posterior samples to those obtained using LALINFERENCE, we reweight the LALINFERENCE posterior distributions by BILBY_PIPE default priors. Appendix C contains the details of this reweighting procedure.

To quantitatively assess the similarity between BILBY and LALINFERENCE posterior samples, we measure their Jensen–Shannon (JS; Lin 1991) divergence. This is a symmetrized extension of the Kullback–Leibler divergence (Kullback & Leibler 1951) that is used to quantify the information gain going between two distributions. The JS divergence is defined to be between 0 bit and 1 bit, where 0 bit represents no additional information going from one distribution to the other (the two distributions are identical) and 1 bit represents maximal divergence. For different sets of samples drawn from the same Gaussian distribution, we find JS divergence values of $\lesssim 0.0010$ bit while the number of samples $N \gtrsim 2000$, and JS divergence values of $\lesssim 0.0004$ bit when $N \gtrsim 5000$. To compare BILBY and LALINFERENCE results, we use $N = \min(N_{\text{LI}}, 10000)$, where N_{LI} is the number of samples left in the LALINFERENCE posterior after the reweighting procedure.

Our goal is to use the JS divergence as a quantitative indicator that the BILBY GWTC-1 samples are in agreement with those produced by LALINFERENCE. To investigate the typical distributions of JS divergence values due to sampling error, we calculated JS values for posteriors from two distinct LALINFERENCE runs on GW150914 with identical configurations. Bootstrapping was used to generate 100 posterior realizations from each run, which were used to obtain a distribution of JS divergences for each of the binary parameters included in the public LALINFERENCE GWTC-1 posterior sample release. Across different parameters, we typically found mean values of 0.0007 bit, with a maximum of 0.0015 bit. As such, we determined the

Figure 10. Joint posterior distributions for parameters of GW170817, comparing PBILBY posteriors in orange and LALINFERENCE posteriors in blue. Left: Posterior probability distributions for tidal parameters $\tilde{\Lambda}$ ($\text{JS}_{\tilde{\Lambda}} = 0.0019$ bit) and $\delta\tilde{\Lambda}$ ($\text{JS}_{\delta\tilde{\Lambda}} = 0.0008$ bit). Right: Posterior probability distributions for inclination angle θ_{JN} ($\text{JS}_{\theta_{\text{JN}}} = 0.0009$ bit) and luminosity distance d_L ($\text{JS}_{d_L} = 0.0008$ bit).

following naive criteria for evaluating the JS divergence values when comparing the BILBY and LALINFERENCE GWTC-1 posteriors. For a JS divergence value less than 0.0015 bit, we conclude the samples are, to within statistical uncertainties, drawn from the same distribution, and values larger than 0.0015 bit require manual inspection.

In Table 1, we list the maximum JS divergence for the model parameters for each event. Of these, six pass our naive criterion described above. For the remaining events, we manually inspect the posterior distributions to look for discrepancies. The parameter with the largest JS divergence value across all BBH events is the right ascension, α . Events with large sky areas, such as GW170729, suffer from large deviations between the BILBY and LALINFERENCE posteriors in the sky position parameters. The sky position was fixed to the location of the EM counterpart for GW170817. We show the difference between the BILBY and LALINFERENCE posterior cumulative density functions (CDFs) for α in Fig. 6 and for the luminosity distance d_L , which passes the naive criterion on the JS divergence for all events, in Fig. 7. For GW170818, α has the largest JS divergence value (0.006 bit) despite the fact that the BILBY and LALINFERENCE CDFs match at the 2σ level. This is because the distribution is approximated using a kernel density estimate (KDE) in order to compute the JS divergence, and the posterior for this particular event has a sharp drop-off, which is difficult to model faithfully using the KDE.

Upon manual inspection, we find that the posteriors with JS divergence values up to ~ 0.002 bit are consistent between the LALINFERENCE and BILBY samples. The remaining parameters with significant deviations between the two samplers are the sky position parameters for GW170729. Investigations into the source of these discrepancies are ongoing. The differences between the BILBY and LALINFERENCE CDFs for all events and all parameters are shown in Appendix D. A similar comparison was made in Abbott et al. (2019f) analysing the posterior distributions obtained using two different waveform approximants for each event. The maximum difference between the posteriors assuming the two different waveform models in that work is typically ~ 0.02 bit, an order of magnitude larger than the differences here.

As another way to visualize the differences between the BILBY and LALINFERENCE samples, in Fig. 8, we compare the 90 per cent

credible areas of the two posteriors on the source-frame primary mass m_1^{source} and secondary mass m_2^{source} for all GWTC-1 events. As indicated by the low JS divergence values for the mass parameters, the two samplers produce posteriors on these parameters that agree within expected statistical fluctuations.

We compare BILBY posteriors on source-frame chirp mass $\mathcal{M}^{\text{source}}$ and luminosity distance d_L for the first observed gravitational-wave event, GW150914 (Abbott et al. 2016a), in Fig. 9. The LALINFERENCE distance posterior here matches the BILBY posterior more closely than was demonstrated in fig. 2 of Ashton et al. (2019). This is due to an issue in the application of the time-domain window being fixed in LALINFERENCE, which had affected the distance posterior (Talbot 2020).

For the first observed binary neutron-star merger event, GW170817, we compare the BILBY posterior distributions on tidal parameters $\tilde{\Lambda}$ and $\delta\tilde{\Lambda}$, as well as θ_{JN} and d_L , to those obtained using LALINFERENCE in Fig. 10. The maximum JS divergence for this event is $\text{JS}_q = 0.0017$ bit. Additional posterior probability plots for all parameters of all 11 CBC events can be found within the online resources that accompany this paper (Romero-Shaw et al. 2020b).

Based on these results, we conclude that BILBY and LALINFERENCE produce statistically indistinguishable results for all parameters and all events reported in GWTC-1 with the exception of the sky area for GW170729 and GW151226. We emphasize that the differences in the CDFs for these parameters are still small compared to other sources of error such as waveform systematics (Abbott et al. 2019f) and uncertainty in the power spectral density (Biscoveanu et al. 2020a). We provide PESUMMARY comparison pages between BILBY and reweighted LALINFERENCE posteriors for all GWTC-1 events online.¹²

5 SUMMARY

BILBY is a modern and versatile Bayesian inference library, and has been primed for analysis of gravitational-wave observations. BILBY performs reliably, producing accurate and unbiased parameter estimation results when analysing simulated signals. We validate

¹²bilby-gwtc1.github.io.

BILBY results for GWTC-1 using the JS divergence statistic between posterior distributions obtained using BILBY and the previously published LALINFERENCE results, finding a maximum JS value of $JS_\alpha = 0.0026$ bit for GW170729. The similarity between the two results indicates that both the BILBY samples obtained with DYNESTY and the LALINFERENCE samples are well-converged, and efforts to further validate these results using alternative samplers within BILBY are ongoing. Posterior probability distributions generated by BILBY and LALINFERENCE, when run on the same GWTC-1 data and using identical analysis settings, are consistent to the level of sampling noise. The BILBY posterior samples for events in GWTC-1 are available online (Romero-Shaw et al. 2020c). We conclude that BILBY is well-suited to meet the challenges of gravitational-wave parameter estimation in the era of frequent detections.

ACKNOWLEDGEMENTS

This work is supported through Australian Research Council (ARC) Centre of Excellence CE170100004. PDL is supported through ARC Future Fellowship FT160100112 and ARC Discovery Project DP180103155. ET is supported through ARC Future Fellowship FT150100281 and CE170100004. This work is partially supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MEST) (No. 2019R1A2C2006787). NB acknowledges Inspire division, DST, Government of India for the fellowship support. This work is partially supported by the National Science Foundation under Grant No. PHY-1912648. SB, C-JH, and CT acknowledge support of the National Science Foundation, and the LIGO Laboratory. SB is also supported by the Paul and Daisy Soros Fellowship for New Americans and the NSF Graduate Research Fellowship under Grant No. DGE-1122374. This work was partially supported by European Union FEDER funds, the Spanish Ministry of Science and Innovation and the Spanish Agencia Estatal de Investigación grants FPA2016-76821-P and PID2019-106416GB-I00/AEI/10.13039/501100011033, the Comunitat Autònoma de les Illes Balears through the Direcció General de Política Universitària i Recerca with funds from the Tourist Stay Tax Law ITS 2017-006 (PRD2018/24), the Vicepresidència i Conselleria d'Innovació, Recerca i Turisme, Conselleria d'Educació, i Universitats del Govern de les Illes Balears and Fons Social Europeu. MC acknowledges funding from the European Union's Horizon 2020 research and innovation programme, under the Marie Skłodowska-Curie grant agreement No. 751492. D.K. is supported by the Spanish Ministerio de Ciencia, Innovación y Universidades (ref. BEAGAL 18/00148) and cofinanced by the Universitat de les Illes Balears. This work used BILBY = v0.6.9, BILBYPIPE = v0.3.12, DYNESTY = v1.0.1, LALSUITE = v6.49, and PESUMMARY = v0.5.6.

This research has made use of data, software and/or web tools obtained from the Gravitational Wave Open Science Center (Abbott et al. 2019e), a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. Computing was performed on the OzSTAR Australian national facility at Swinburne University of Technology, which receives funding in part from the Astronomy National Collaborative Research Infrastructure Strategy (NCRIS) allocation provided by the Australian Government, LIGO Laboratory computing clusters at California Institute of Technology and LIGO Hanford Observatory supported by National Science Foundation Grants PHY-0757058 and PHY-0823459, and the Quest computing cluster, which is jointly supported by the Office of the Provost, the Office for Research and Northwestern University Information Technology, and funded by the National Science Foundation under Grant No. PHY-1726951. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology

with funding from the National Science Foundation and operates under cooperative agreement PHY-1764464. Virgo is funded by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale della Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by Polish and Hungarian institutes.

DATA AVAILABILITY STATEMENT

We analyse publicly available data (Abbott et al. 2019e), and make use of publicly available PSDs (Abbott et al. 2019b) and calibration envelopes (Abbott et al. 2019c). We compare our results against publicly available posterior samples (Abbott et al. 2018a). We make our own results publicly accessible online (Romero-Shaw et al. 2020c).

REFERENCES

- Aasi J. et al., 2013, *Phys. Rev. D*, 88, 062001
Aasi J. et al., 2015, *Class. Quantum Gravity*, 32, 115012
Abbott B. P. et al., 2016a, *Phys. Rev. Lett.*, 116, 061102
Abbott B. P. et al., 2016b, *Phys. Rev. Lett.*, 116, 221101
Abbott B. P. et al., 2016c, *Phys. Rev. Lett.*, 116, 241102
Abbott B. P. et al., 2016d, *ApJ*, 818, L22
Abbott B. P. et al., 2017a, *Class. Quantum Gravity*, 34, 104002
Abbott B. P. et al., 2017b, *Phys. Rev. Lett.*, 118, 221101
Abbott B. P. et al., 2017c, *Phys. Rev. Lett.*, 119, 161101
Abbott B. P. et al., 2017d, *Nature*, 551, 85
Abbott B. P. et al., 2017e, *ApJ*, 848, L12
Abbott B. P. et al., 2017f, *ApJ*, 848, L13
Abbott B. P. et al., 2017g, *ApJ*, 850, L39
Abbott B. P. et al., 2017h, *ApJ*, 850, L40
Abbott B. P. et al., 2018a, Parameter Estimation Sample Release for GWTC-1, <https://dcc.ligo.org/LIGO-P1800370/public>
Abbott B. P. et al., 2018b, *Living Rev. Relativ.*, 21, 3
Abbott B. P. et al., 2018c, *Phys. Rev. Lett.*, 121, 161101
Abbott B. P. et al., 2019c, *Class. Quantum Grav.*, 37, 055002
Abbott B. P. et al., 2019b, Power Spectral Densities (PSD) Release for GWTC-1, <https://dcc.ligo.org/LIGO-P1900011/public>
Abbott B. P. et al., 2019c, Calibration Uncertainty Envelope Release for GWTC-1, <https://dcc.ligo.org/LIGO-P1900040/public>
Abbott B. P. et al., 2019d, preprint (arXiv:1908.06060)
Abbott R. et al., 2019e, preprint (arXiv:1912.11716)
Abbott B. P. et al., 2019f, *Phys. Rev. X*, 9, 031040
Abbott B. P. et al., 2019g, *Phys. Rev. X*, 9, 011001
Abbott B. P. et al., 2019h, *Phys. Rev. D*, 100, 104036
Abbott B. P. et al., 2019i, *ApJ*, 882, L24
Abbott B. P. et al., 2020a, *Phys. Rev. D*, 102, 043015
Abbott B. P. et al., 2020b, *Class. Quantum Gravity*, 37, 045006
Abbott R. et al., 2020c, *Phys. Rev. Lett.*, 125, 101102
Abbott B. P. et al., 2020d, *ApJ*, 892, L3
Acernese F. et al., 2015, *Class. Quantum Gravity*, 32, 024001
Ade P. A. et al., 2016, *A&A*, 594, A13
Ajith P. et al., 2011, *Phys. Rev. Lett.*, 106, 241101
Allen B., Anderson W. G., Brady P. R., Brown D. A., Creighton J. D. E., 2012, *Phys. Rev. D*, 85, 122006
Ashton G., Khan S., 2020, *Phys. Rev. D*, 101, 064037
Ashton G. et al., 2019, *ApJS*, 241, 27
Babak S. et al., 2008, *Class. Quantum Gravity*, 25, 184026
Babak S. et al., 2010, *Class. Quantum Gravity*, 27, 084009
Barrett J. W., Gaebel S. M., Neijssel C. J., Vigna-Gómez A., Stevenson S., Berry C. P. L., Farr W. M., Mandel I., 2018, *MNRAS*, 477, 4685
Bavera S. S. et al., 2020, *A&A*, 635, A97
Bayes T., 1763, *Phil. Trans. R. Soc.*, 53, 370
Baylor A., Smith R., Chase E., 2019, IMRPhenomPv2_NRTidal_GW190425_narrow_Mc [Data set]. Zenodo
Belczynski K. et al., 2018, preprint (arXiv:1812.10065)
Berry C. P. L. et al., 2015, *ApJ*, 804, 114
Biscoveanu S., Vitale S., Haster C.-J., 2019, *ApJ*, 884, L32

- Biscoveanu S., Haster C.-J., Vitale S., Davies J., 2020a, *Phys. Rev. D*, 102, 023008
- Biscoveanu S., Thrane E., Vitale S., 2020b, *ApJ*, 893, 38
- Biwer C. M., Capano C. D., De S., Cabero M., Brown D. A., Nitz A. H., Raymond V., 2019, *PASP*, 131, 024503
- Blackman J. et al., 2017, *Phys. Rev. D*, 96, 024058
- Blanchet L., Damour T., Iyer B. R., Will C. M., Wiseman A. G., 1995, *Phys. Rev. Lett.*, 74, 3515
- Bohé A., Hannam M., Husa S., Ohme F., Pürrer M., Schmidt P., 2016, Technical Report LIGO-T1500602, PhenomPv2 – Technical Notes for LAL Implementation, <https://dcc.ligo.org/LIGO-T1500602/public>
- Bohé A. et al., 2017, *Phys. Rev. D*, 95, 044028
- Buchner J., 2016, *Stat. Comput.*, 26, 383
- Buchner J., 2019, *PASP*, 131, 108005
- Buchner J. et al., 2014, *A&A*, 564, A125
- Cahillane C. et al., 2017, *Phys. Rev. D*, 96, 102001
- Cantiello M. et al., 2018, *ApJ*, 854, L31
- Chatziioannou K., Haster C.-J., Littenberg T. B., Farr W. M., Ghonge S., Millhouse M., Clark J. A., Cornish N., 2019, *Phys. Rev. D*, 100, 104004
- Chen H.-Y., Fishbach M., Holz D. E., 2018, *Nature*, 562, 545
- Chornock R. et al., 2017, *ApJ*, 848, L19
- Christensen N., Meyer R., 1998, *Phys. Rev. D*, 58, 082001
- Christensen N., Meyer R., 2001, *Phys. Rev. D*, 64, 022001
- Cook S. R., Gelman A., Rubin D. B., 2006, *J. Comput. Graph. Stat.*, 15, 675
- Cornish N. J., 2010, preprint ([arXiv:1007.4820](https://arxiv.org/abs/1007.4820))
- Cornish N. J., Littenberg T. B., 2015, *Class. Quantum Gravity*, 32, 135012
- Cornish N. J., Shuman K., 2020, *Phys. Rev. D*, 101, 124008
- Coughlin M. W., Dietrich T., 2019, *Phys. Rev. D*, 100, 043011
- Del Pozzo W., Berry C. P. L., Ghosh A., Haines T. S. F., Singer L. P., Vecchio A., 2018, *MNRAS*, 479, 601
- Dhawan S., Bulla M., Goobar A., Sagués Carracedo A., Setzer C. N., 2020, *ApJ*, 888, 67
- Dietrich T., Samajdar A., Khan S., Johnson-McDaniel N. K., Dudi R., Tichy W., 2019, *Phys. Rev. D*, 100, 044003
- Ellis J., van Haasteren R., 2017, jellis18/PTMCMCSampler: Official Release (Version 1.0.0). Zenodo
- Essick R., Landry P., Holz D. E., 2020, *Phys. Rev. D*, 101, 063007
- Farah W. et al., 2019, *MNRAS*, 488, 2989
- Farr B., Farr W. M., 2015
- Farr W. M., Farr B., Littenberg T., 2014, Technical Report LIGO-T1400682, Modelling Calibration Errors in CBC Waveforms, <https://dcc.ligo.org/LIGO-T1400682/public>
- Farr W. M., Mandel I., Aldridge C., Stroud K., 2014, Technical Report LIGO-T1400460, Marginalisation of the Time and Phase Parameters in CBC Parameter Estimation, <https://dcc.ligo.org/public/0114/T1400460/002/margtime.pdf>
- Favata M., 2014, *Phys. Rev. Lett.*, 112, 101101
- Feroz F., Hobson M. P., 2008, *MNRAS*, 384, 449
- Feroz F., Hobson M. P., Bridges M., 2009, *MNRAS*, 398, 1601
- Feroz F., Hobson M. P., Cameron E., Pettitt A. N., 2019, *Open J. Astrophys.*, 2, 10
- Finn L. S., 1992, *Phys. Rev. D*, 46, 5236
- Finn L. S., Chernoff D. F., 1993, *Phys. Rev. D*, 47, 2198
- Fishbach M., Holz D. E., 2017, *ApJ*, 851, L25
- Flanagan É. É., Hinderer T., 2008, *Phys. Rev. D*, 77, 021502
- Fong W. et al., 2019, *ApJ*, 883, L1
- Gabbard H., Messenger C., Heng I. S., Tonolini F., Murray-Smith R., 2019, preprint ([arXiv:1909.06296](https://arxiv.org/abs/1909.06296))
- Galaudage S., Talbot C., Thrane E., 2019, preprint ([arXiv:1912.09708](https://arxiv.org/abs/1912.09708))
- George D., Huerta E. A., 2018, *Phys. Lett. B*, 778, 64
- Goncharov B., Zhu X.-J., Thrane E., 2020, *MNRAS*, 497, 3264
- Górski K. M. et al., 1999, in Banday A. J., Sheth R. K., da Costa L. N., eds, *Evolution of Large Scale Structure : From Recombination to Garching: proceedings of the MPA-ESO Cosmology Conference: Garching, Germany*, p. 37
- Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, 622, 759
- Handley W. J., Hobson M. P., Lasenby A. N., 2015a, *MNRAS*, 450, L61
- Handley W. J., Hobson M. P., Lasenby A. N., 2015b, *MNRAS*, 453, 4384
- Hannam M., Schmidt P., Bohé A., Haegel L., Husa S., Ohme F., Pratten G., Pürrer M., 2014, *Phys. Rev. Lett.*, 113, 151101
- Hastings W. K., 1970, *Biometrika*, 57, 97
- Hernandez Vivanco F., Smith R., Thrane E., Lasky P. D., 2019a, *Phys. Rev. D*, 100, 043023
- Hernandez Vivanco F., Smith R., Thrane E., Lasky P. D., Talbot C., Raymond V., 2019b, *Phys. Rev. D*, 100, 103009
- Hogg D. W., Foreman-Mackey D., 2018, *ApJS*, 236, 11
- Hotokezaka K., Nakar E., Gottlieb O., Nissanke S., Masuda K., Hallinan G., Mooley K. P., Deller A. T., 2019, *Nat. Astron.*, 3, 940
- Hoy C., Raymond V., 2020, preprint ([arXiv:2006.06639](https://arxiv.org/abs/2006.06639))
- Hübner M., Talbot C., Lasky P. D., Thrane E., 2020, *Phys. Rev. D*, 101, 023011
- Isi M., Giesler M., Farr W. M., Scheel M. A., Teukolsky S. A., 2019, *Phys. Rev. Lett.*, 123, 111102
- Kasliwal M. M. et al., 2019, *MNRAS*, L14
- Keitel D., 2019, *Res. Notes Am. Astron. Soc.*, 3, 46
- Khan S., Husa S., Hannam M., Ohme F., Pürrer M., Forteza X. J., Bohé A., 2016, *Phys. Rev. D*, 93, 044007
- Kimball C., Talbot C., Berry C. P. L., Carney M., Zevin M., Thrane E., Kalogera V., 2020, *ApJ*, 900, 177
- Kolmogorov A. N., 1933, *G. Ist. Ital. Attuari.*, 4, 83
- Krolak A., Schutz B. F., 1987, *Gen. Relativ. Gravit.*, 19, 1163
- Kullback S., Leibler R. A., 1951, *Ann. Math. Statist.*, 22, 79
- Lackey B. D., Pürrer M., Taracchini A., Marsat S., 2019, *Phys. Rev. D*, 100, 024002
- Lange J., O'Shaughnessy R., Rizzo M., 2018, preprint ([arXiv:1805.10457](https://arxiv.org/abs/1805.10457))
- Lentati L., Alexander P., Hobson M. P., Feroz F., van Haasteren R., Lee K. J., Shannon R. M., 2014, *MNRAS*, 437, 3004
- LIGO Scientific Collaboration, 2018, LIGO Algorithm Library - LALSuite, Free Software (GPL), <https://doi.org/10.7935/GT1W-FZ16>
- Lin J., 1991, *IEEE Trans. Inf. Theory*, 37, 145
- Littenberg T. B., Cornish N. J., 2015, *Phys. Rev. D*, 91, 084034
- Lower M. E., Thrane E., Lasky P. D., Smith R., 2018, *Phys. Rev. D*, 98, 083028
- Macleod D. et al., 2018, gwpy/gwpy: 0.12.0 (Version v0.12.0). Zenodo
- Margutti R. et al., 2018, *ApJ*, 856, L18
- Marsat S., Baker J. G., Dal Canton T., 2020, preprint ([arXiv:2003.00357](https://arxiv.org/abs/2003.00357))
- Metropolis N., Rosenbluth A. W., Rosenbluth M. N., Teller A. H., Teller E., 1953, *J. Chem. Phys.*, 21, 1087
- Mooley K. P. et al., 2018, *Nature*, 561, 355
- Most E. R., Weih L. R., Rezzolla L., Schaffner-Bielich J., 2018, *Phys. Rev. Lett.*, 120, 261103
- Nagar A. et al., 2018, *Phys. Rev. D*, 98, 104052
- Ossokine S. et al., 2020, *Phys. Rev. D*, 102, 44055
- Pankow C., Brady P., Ochsner E., O'Shaughnessy R., 2015, *Phys. Rev. D*, 92, 023002
- Pankow C. et al., 2018, *Phys. Rev. D*, 98, 084016
- Payne E., Talbot C., Thrane E., 2019, *Phys. Rev. D*, 100, 123017
- Poisson E., Will C. M., 1995, *Phys. Rev. D*, 52, 848
- Powell J., 2018, *Class. Quantum Gravity*, 35, 155017
- Powell J., Müller B., 2019, *MNRAS*, 487, 1178
- Price-Whelan A. M. et al., 2018, *AJ*, 156, 18
- Ramos-Buades A., Husa S., Pratten G., Estellés H., García-Quirós C., Mateu-Lucena M., Colleoni M., Jaume R., 2020, *Phys. Rev. D*, 101, 083015
- Robitaille T. P. et al., 2013, *A&A*, 558, A33
- Romano J. D., Cornish N. J., 2017, *Living Rev. Relativ.*, 20, 2
- Romero-Shaw I. M., Lasky P. D., Thrane E., 2019, *MNRAS*, 5210, 490
- Romero-Shaw I. M., Farrow N., Stevenson S., Thrane E., Zhu X.-J., 2020a, *MNRAS*
- Romero-Shaw I. et al., 2020b, IsobelMarguarethe/Bilby-GWTC-1-Analysis-and-Verification: v2.0 (Version v2.0). Zenodo
- Romero-Shaw I. et al., 2020c, <https://dcc.ligo.org/LIGO-P2000193/public>
- Röver C., Meyer R., Christensen N., 2006, *Class. Quantum Gravity*, 23, 4895
- Röver C., Meyer R., Christensen N., 2007, *Phys. Rev. D*, 75, 062004
- Santamaria L. et al., 2010, *Phys. Rev. D*, 82, 064016
- Sarin N., Lasky P. D., Ashton G., 2020, *Phys. Rev. D*, 101, 063021
- Schmidt P., Hannam M., Husa S., 2012, *Phys. Rev. D*, 86, 104063
- Schmidt P., Ohme F., Hannam M., 2015, *Phys. Rev. D*, 91, 024043

- Sidery T. et al., 2014, *Phys. Rev. D*, 89, 084060
- Singer L. P., Price L. R., 2016, *Phys. Rev. D*, 93, 024013
- Singer L. P. et al., 2014, *ApJ*, 795, 105
- Singer L. P. et al., 2016, *ApJ*, 829, L15
- Singer L. P. et al., 2020, GWCElery. <https://git.ligo.org/emfollow/gwcelery>
- Skilling J., 2006, *Bayesian Anal.*, 1, 833
- Smirnov N., 1948, *Ann. Math. Statist.*, 19, 279
- Smith R., Field S. E., Blackburn K., Haster C.-J., Pürrer M., Raymond V., Schmidt P., 2016, *Phys. Rev. D*, 94, 44031
- Smith R., Ashton G., Vajpeyi A., Talbot C., 2020, *MNRAS*, 498, 4492
- Sokal A. D., 1996, Monte Carlo Methods for the Self-Avoiding Walk, *Nuclear Physics B - Proceedings Supplements*, 47, 172
- Speagle J. S., 2020, *MNRAS*, 493, 3132
- Stevenson S., Ohme F., Fairhurst S., 2015, *ApJ*, 810, 58
- Suwa H., Todo S., 2010, *Phys. Rev. Lett.*, 105, 120603
- Talbot C., 2020, PhD thesis, Monash University
- Talbot C., Smith R., Thrane E., Poole G. B., 2019, *Phys. Rev. D*, 100, 043030
- Talts S., Betancourt M., Simpson D., Vehtari A., Gelman A., 2018, preprint ([arXiv:1804.06788](https://arxiv.org/abs/1804.06788))
- Tanvir N. R. et al., 2017, *ApJ*, 848, L27
- Thrane E., Talbot C., 2019, *Publ. Astron. Soc. Aust.*, 36, e010
- van der Sluys M., Raymond V., Mandel I., Röver C., Christensen N., Kalogera V., Meyer R., Vecchio A., 2008a, *Class. Quantum Gravity*, 25, 184011
- van der Sluys M. V. et al., 2008b, *ApJ*, 688, L61
- Veitch J., Del Pozzo W., 2013, Technical Report LIGO-T1300326, Analytic Marginalisation of Phase Parameter, <https://dcc.ligo.org/LIGO-T1300326/public>
- Veitch J., Vecchio A., 2008, *Phys. Rev. D*, 78, 022001
- Veitch J., Vecchio A., 2010, *Phys. Rev. D*, 81, 062003
- Veitch J. et al., 2015, *Phys. Rev. D*, 91, 042003
- Veitch J., Pozzo W. D., Cody Pitkin M., 2017, johnveitch/cpnest: Minor optimisation (Version 0.1.4). Zenodo
- Viets A. D. et al., 2018, *Class. Quantum Gravity*, 35, 095015
- Vigeland S. J., Vallisneri M., 2014, *MNRAS*, 440, 1446
- Vitale S., Del Pozzo W., Li T. G. F., Van Den Broeck C., Mandel I., Aylott B., Veitch J., 2012, *Phys. Rev. D*, 85, 064034
- Wade L., Creighton J. D. E., Ochsner E., Lackey B. D., Farr B. F., Littenberg T. B., Raymond V., 2014, *Phys. Rev. D*, 89, 103012
- Wang S., Zhao Z.-C., 2020, preprint ([arXiv:2002.00396](https://arxiv.org/abs/2002.00396))
- Watson D. et al., 2019, *Nature*, 574, 497
- Yunes N., Siemens X., 2013, *Living Rev. Relativ.*, 16, 9
- Yunes N., Yagi K., Pretorius F., 2016, *Phys. Rev. D*, 94, 084002
- Zackay B., Dai L., Venumadhav T., 2018, preprint ([arXiv:1806.08792](https://arxiv.org/abs/1806.08792))
- Zevin M., Pankow C., Rodriguez C. L., Sampson L., Chase E., Kalogera V., Rasio F. A., 2017, *ApJ*, 846, 82
- Zevin M., Berry C. P. L., Coughlin S., Chatziioannou K., Vitale S., 2020, *ApJ*, 899, L17
- Zhao Z.-C., Lin H.-N., Chang Z., 2019, *Chin. Phys. C*, 43, 075102

APPENDIX A: ADDITIONAL BILBY VALIDATION TESTS

In addition to the tests described in the main body of the paper, we performed several additional validation tests that are standard benchmarks for stochastic sampling codes.

A1 Prior sampling

The initial distribution of samples drawn from the prior must faithfully represent the shape of the prior function. In addition to being used for review, the prior sampling test also forms part of BILBY's unit test suite. Prior samples can be obtained using BILBY via two different methods. The first is to use the `sample` method of each `Prior` object, which generates samples by rescaling from a unit cube. The second is to run the sampler with a null likelihood using the `ZeroLikelihood` object so that the returned posterior samples

actually reflect the prior. To test the consistency of the two methods, we generate prior samples via both methods for a standard 15D binary black hole signal injected into simulated Gaussian noise. We perform a Kolmogorov–Smirnov test (Kolmogorov 1933; Smirnov 1948) to evaluate the similarity of the two sets of samples, calculating a p -value for each parameter, which quantifies the probability that the two sets of samples are drawn from identical distributions. A combined p -value is then computed, representing the probability that the ensemble of individual-parameter p -values is drawn from a unit uniform distribution. We consider the test to pass if this combined p -value is greater than 0.01. For a representative run with the latest version of BILBY, we obtain a combined p -value of 0.017.

A2 15D Gaussian

Sampling an analytically known likelihood distribution is an important test to verify that we can recover the correct posterior. For this test, we choose the SCIPY implementation of a multivariate normal distribution (`scipy.stats.multivariate.normal`) as our likelihood. We choose the distribution to be 15D since this reflects the typical number of dimensions we encounter in binary black hole problems. We set the means of all parameters to be zero, and choose a covariance matrix COV_{ij} with standard deviations for each of the parameters ranging between 0.15 and 0.25 to match past tests done with LALINFERENCE. Using the BILBY default sampler settings for a 15D problem, we test if we correctly recover the posterior distribution by drawing samples from this 15D likelihood and comparing the obtained means and standard deviations to the true values. Additionally, we verify that we recover the expected evidence within the estimated error. Since the likelihood distribution is normalized and we use uniform priors for each parameter in the range $[-5, 5]$, the evidence can be approximated by the prior volume, since the standard deviations are small enough that the value of the likelihood evaluated at the edges of the prior is negligible

$$\ln \mathcal{Z} \approx -\ln X, \quad (\text{A1})$$

where X is the prior volume. In Fig. A1 on the left-hand side we find the measured standard deviations and the evidence to be in broad agreement with analytical expectations. While the evidence errors quoted by DYNESTY are not truly Gaussian, the 1σ credible interval is consistent with covering the true evidence 68 per cent of the time if one uses more than 1000 live points. Additionally, the overshoot at high values of the credible interval indicates that there are fewer outliers than we would for a Gaussian distribution. The right-hand side of Fig. A1 demonstrates that the width of the posterior distribution is correctly recovered. We have thus shown that the DYNESTY implementation in BILBY has no significant issues in recovering the shape of posterior distributions and the correct evidence for this fundamental problem.

We performed the same test using a bimodal Gaussian distribution, with means separated by eight standard deviations in each dimension. While it is more difficult to correctly sample a degenerate likelihood surface, we still find 1000 live points sufficient to reasonably recover the evidence. Individual runs of the bimodal likelihood may produce a biased set posterior samples in favour of one of the modes over the other, which is why multiple runs should be combined. We verified that none of the modes is preferred if we use all 100 runs. Thus, there are also no substantial issues that arise in sampling multimodal distributions with BILBY.

Figure A1. Left: Illustration of the frequency with which the true evidence is within a given credible interval for the unimodal Gaussian-shaped likelihood. The legend shows how many live points are used to produce the individual curves. For lower number of live points, systematic errors in the evidence estimation cause significant underestimates of the error. Starting at 1024 live points, the evidence error reasonably reflects the true uncertainty. The grey band shows the 90 per cent confidence interval. Right: Residuals of the true width of the analytical likelihood minus the average recovered one for 1024 live points in each dimension based on 100 independent runs. The error bars show the 90 per cent confidence interval of the average mean of the distribution. There is a small $\mathcal{O}(0.1$ per cent) systematic bias to underestimate the width, i.e. the parameter is on average slightly overconstrained. However, this bias is negligibly small compared to stochastic sampling uncertainties for individual runs.

Table A1. Our injected and recovered values for the two fiducial event analyses. Recovered median values are quoted with the symmetric 90 per cent credible interval around the median.

Parameter	BBH		BNS	
	Inject	Recover	Inject	Recover
$\mathcal{M}/M_{\odot}$	15.53	$15.4^{+0.3}_{-0.4}$	1.486	$1.486^{+0.0001}_{-0.0001}$
q	0.52	$0.7^{+0.3}_{-0.4}$	0.9	$0.9^{+0.1}_{-0.2}$
a_1	0.65	$0.6^{+0.3}_{-0.5}$	0.04	$0.02^{+0.02}_{-0.02}$
a_2	0.65	$0.5^{+0.4}_{-0.4}$	0.01	$0.02^{+0.02}_{-0.02}$
θ_1	1.24	$1.1^{+0.8}_{-0.6}$	1.03	$1.5^{+1.0}_{-0.9}$
θ_2	0.80	$1.3^{+1.1}_{-0.9}$	2.17	$1.6^{+1.0}_{-1.0}$
ϕ_{12}	1.5	$3.1^{+2.9}_{-2.8}$	5.10	$3.2^{+2.8}_{-2.9}$
ϕ_{JL}	3.01	$3.2^{+2.8}_{-2.9}$	2.52	$3.1^{+2.9}_{-2.8}$
d_L/Mpc	614	1018^{+1147}_{-623}	100	86^{+17}_{-26}
δ	1.00	$0.7^{+0.4}_{-1.6}$	0.2	$0.3^{+0.1}_{-0.1}$
α	2.00	$4.6^{+1.0}_{-2.7}$	3.95	$3.9^{+0.1}_{-0.1}$
θ_{JN}	1.65	$1.8^{+1.0}_{-0.8}$	0.25	$0.6^{+0.7}_{-0.4}$
ψ	1.50	$1.6^{+1.4}_{-1.4}$	2.70	$1.5^{+1.5}_{-1.4}$
ϕ	2.00	$3.1^{+2.8}_{-2.8}$	3.69	$3.1^{+2.8}_{-2.8}$
t_{geo}/s	0.04	$0.04^{+0.00}_{-0.02}$	-0.01	$-0.01^{+0.00}_{-0.00}$
Λ_1	—	—	1500	752^{+915}_{-657}
Λ_2	—	—	750	1437^{+1294}_{-1216}

A3 Fiducial event simulations

We analyse two fiducial simulated signals; one binary black hole merger, and one binary neutron star merger with tides. We use a LIGO Hanford–Livingston detector network and add the simulated signals into design sensitivity Gaussian noise. For the binary black hole, we use the IMRPHENOMPv2 waveform and the default 4 s prior described in Table B1. For the binary neutron star, we use the ROQ implementation of the IMRPhenomPv2_NRTidalv2 waveform (Baylor, Smith & Chase 2019) with the 128 s tidal low-spin prior. The binary black hole and neutron star systems have network optimal

SNRs of 8.8 and 27.9, respectively.¹³ In Table A1, we show the true values along with the recovered median and 90 per cent credible interval values for each parameter. Nearly all the true parameter values for both systems are recovered within the 90 per cent credible interval, and those that are not are consistent with deviations due to the Gaussian noise realization. Full corner plots for both simulated signals are available online (Romero-Shaw et al. 2020b).

APPENDIX B: RUN SETTING DETAILS

B1 Sampler settings

The default sampler used by BILBY is DYNESTY (Speagle 2020), an off-the-shelf nested sampling (Skilling 2006) package. The first step in nested sampling is to draw N random live points from the prior. At each iteration, the lowest likelihood sample from the initial N points is discarded in favour of a higher likelihood point, again randomly chosen from the prior. After every step, the actively sampled region of the prior shrinks to the volume contained by the hyperplane of constant minimum likelihood for the current population of live points. When the live domain has reduced sufficiently, it becomes inefficient to select higher likelihood points uniformly from the restricted prior space.

After the uniform sampling becomes sufficiently inefficient, new points are selected by randomly walking using a custom Markov-chain Monte Carlo algorithm starting from the sample being replaced. The transition probability is determined by the distribution of the set of current live points. The number of steps taken in the chain is determined such that the length of the chain is at least some multiple n_{act} of the autocorrelation length of the chain (Sokal 1994). For the analysis in this paper, we require $n_{\text{act}} = 10$. A Markov-chain Monte Carlo walker algorithm then takes at least n steps to draw a

¹³The binary black hole analysis was performed using BILBY version 0.6.3, while the neutron star analysis used BILBY 1.0.0. The default Advanced LIGO design PSD changed between these two versions of BILBY to reflect the updated detector sensitivity predictions (Abbott et al. 2018b). Parameter estimation is performed using DYNESTY with the default settings.

Table B1. Lower and upper limits on chirp mass $\mathcal{M}$, luminosity distance d_L , and dimensionless spin magnitude a_1 , a_2 priors for each of the default prior sets contained in BILBY_PIPE.

Prior	$\mathcal{M}/M_\odot$	d_L/Mpc	a_1, a_2
High mass	25–175	100–7000	0–0.99
4 s	12.299703–45	100–5000	0–0.88
8 s	7.932707–14.759644	100–5000	0–0.8
16 s	5.141979–9.519249	100–4000	0–0.8
32 s	3.346569–6.170374	100–3000	0–0.8
64 s	2.184345–4.015883	20–2000	0–0.8
128 s	1.420599–2.602169	1–500	0–0.8
128 s tidal	1.485–1.49	1–300	0–0.89
128 s tidal low spin	1.485–1.49	1–300	0–0.05

new sample from the restricted prior. In order to reduce bottlenecks while using multiprocessing we impose a maximum length of the chain. If no point with a higher likelihood than the original point is found within this number of steps, we return a random point from the prior distribution. Nested sampling is able to well-resolve multimodal distributions, making it useful for exploring complicated parameter spaces. For all events in GWTC-1, we give the sampler $N = 2000$ live points and $n = 100$ steps.

B2 Priors

We sample directly in $\mathcal{M}$ and q to avoid issues associated with sampling extremely thin regions of parameter space, which occurs when sampling in component masses (BILBY and BILBY_PIPE can easily be made to sample in other parameters such as component masses; here, we only discuss default parameters and priors used for analysis of the 11 events in GWTC-1). Our prior on mass ratio is uniform in the range $0.125 \leq q \leq 1.0$, with the lower limit determined due to limitations of the IMRPHENOMPV2 ROQ.

Prior limits used for $\mathcal{M}$, d_L , a_1 , and a_2 are provided in Table B1. The chirp mass prior limits are based on those stated in the ROQ git repository.⁶ We use a luminosity distance prior that is uniform in the source frame, with limits motivated by the scaling of gravitational-wave amplitude with both chirp mass and distance. The uniform-in-source-frame prior, which indicates a uniform distribution of mergers in our Universe (Ade et al. 2016), differs from the d_L^2 power-law prior used in the LALINFERENCE analyses, which indicates a uniform distribution in a Euclidean, non-expanding universe. We use dimensionless component spin priors that are uniform between 0 and an upper limit that is determined by the mass range assumed. For non-tidal waveform models, we use an upper limit that is either 0.8, 0.88, or 0.99. For tidal approximants, both a low-spin and a high-spin prior are available. Our component spin prior upper limits are 0.05 (low-spin) and 0.89 (high-spin) in these cases. The upper limits on spin magnitude are determined by the training range of the ROQ basis (e.g. Smith et al. 2016). For analysis of binary neutron star coalescence signal GW170817, we sample in the dimensionless tidal parameters Λ_1 and Λ_2 , which describe the deformability of the primary and secondary masses. If $\Lambda_i = 0$, the neutron star is non-deformable and thus has no tides. We set our priors on Λ_1 and Λ_2 to be uniform between 0 and 5000 to reflect our ignorance of the neutron star equation of state. The remainder of our priors, which are standard and geometrically motivated, have limits as given in Table B2.

Table B2. Default prior settings for 10 of the 17 parameters studied for CBCs observed with gravitational waves. The settings given in this table are consistent between all default prior sets contained in BILBY_PIPE.

Parameter	Shape	Limits	Boundary
q	Uniform	0.125–1	–
θ_1, θ_2	Sinusoidal	0– π	–
ϕ_{12}, ϕ_{JL}	Uniform	0– 2π	Periodic
θ_{JN}	Sinusoidal	0– π	–
ψ	Uniform	0– π	Periodic
ϕ	Uniform	0– 2π	Periodic
α	Uniform	0– 2π	Periodic
δ	Cosinusoidal	$-\pi/2$ – $\pi/2$	–

Table B3. GPS trigger time and data segment duration used for each event. By default, the data segment is positioned such that there are 2 s of data after the trigger time.

Event	GPS trigger time $t_{\text{trig}} \text{ s}^{-1}$	Data duration $T \text{ s}^{-1}$
GW150914	1126259462.391	8
GW151012	1128678900.400	8
GW151226	1135136350.600	8
GW170104	1167559936.600	4
GW170608	1180922494.500	16
GW170729	1185389807.300	4
GW170809	1186302519.700	4
GW170814	1186741861.500	4
GW170817	1187008882.430	128
GW170818	1187058327.100	4
GW170823	1187529256.500	4

B3 Data

The data segments we use are accessed using the GWpy (Macleod et al. 2018) method `TimeSeries.get(channel_name, start_time, end_time)`. The `start_time` t_{start} and `end_time` t_{end} are defined relative to the `trigger_time` t_{trig} of each event, such that

$$t_{\text{end}} = t_{\text{trig}} + t_{\text{post-trig}}; \quad t_{\text{start}} = t_{\text{end}} - T. \quad (\text{B1})$$

Here, T is the total duration of the data segment and $t_{\text{post-trig}}$ is the post-trigger duration, which is 2 s in BILBY by default. We provide the trigger times and data segment durations for all GWTC-1 events in Table B3. The `channel_name` used to obtain strain data from both the LIGO Hanford and LIGO Livingston detectors is `DCS-CALIB_STRAIN_C02` for all events, with the exception of GW170817, for which we use the `channel_name` of `DCH-CLEAN_STRAIN_C02.T1700406_v3` to obtain glitch-subtracted strain data from LIGO Livingston. We also obtain Virgo data for events that occurred from 2017 July until mid-August (GW170729, GW170809, GW170814, GW170817, and GW170818) using the `channel_name` of `Hrec_hoft_V102Repro2A.16384Hz`.

Strain data are available from the Gravitational Wave Open Science Centre (Abbott et al. 2019e) sampled at both 16384 Hz (the native sampling frequency of advanced LIGO and advanced Virgo) and down-sampled to 4096 Hz. We download the data sampled at 16384 Hz. The LALINFERENCE (LIGO Scientific Collaboration 2018) analysis of binary black holes in Abbott et al. (2019f) was performed with data down-sampled to 2048 Hz using a LAL down-sampling function and integrated to the Nyquist frequency (1024 Hz).

In BILBY_PIPE the user can choose to either not down-sample, down-sample using the same LAL routine as done in LALINFERENCE and BAYESWAVE (Cornish & Littenberg 2015), or down-sample

Figure B1. The data and PSD in the LIGO Livingston interferometer at the time of GW170608. In the upper/lower panel we show the data with/without being low-pass filtered and down sampled to 2048 Hz. We can see the effect of the low-pass filter in suppressing the data above ~ 900 Hz. The filtering and down sampling was applied when computing the PSD and so the data on the left better matches the PSD.

using the GWpy method. In general, we recommend users do not down-sample the time domain data, but rather apply cuts directly in the frequency domain. However, since the PSDs used in this analysis were made with BAYESWAVE and the LALINFERENCE analysis we compare with use the LAL down-sampling, we also use this method.

The default method implemented in LAL and used by LALINFERENCE and BAYESWAVE is done in the time domain and consists of two stages. First, the data are low-passed using a 20th-order zero-phase Butterworth filter. The filter is customized such that the power at the low-pass frequency f_c is reduced by a factor of 10. The frequency response of the filter is given by

$$R(f; f_c, n, a_c) = \left[1 + (a_c^{-1/2} - 1) \left(\frac{f}{f_c} \right)^{2n} \right]^{-1}. \quad (\text{B2})$$

The data are then down-sampled by a factor of N by taking every N th sample, this aliases the data. This aliasing means that any signal close to the new Nyquist frequency will be suppressed and aliased

which may introduce a bias in our inference. The final frequency domain strain after downsampling by a factor of N is given by

$$\begin{aligned} \bar{h}(f; f_c, n, a_c) &= h(f)R(f; f_c, n, a_c) \\ &+ \sum_{i=\text{odd}}^N h((i+1)f_c - f)R((i+1)f_c - f; f_c, n, a_c) \\ &+ \sum_{i=\text{even}}^N h(if_c + f)R(if_c + f; f_c, n, a_c). \end{aligned} \quad (\text{B3})$$

Here, $h(f)$ is the frequency-domain data without low-pass filtering or downsampling. Of the events analysed in this work, the lowest mass events (GW151226, GW170608, and GW170817) have frequency content close to or above the down-sampled Nyquist frequency. We expect the bias introduced by this to be small.

In Fig. B1, we show the data containing GW170608 along with the PSD produced by BAYESWAVE with (left) and without (right) downsampling the data to a new sampling rate of 2048 Hz for the LIGO Livingston observatory. On the right we can see the turnover in the data and the PSD close to the new Nyquist frequency 1024 Hz.

APPENDIX C: PRIOR REWEIGHTING

In order to compare posterior samples that are unbiased by differing prior choices, we reweight samples obtained using LALINFERENCE priors π_{LI} by BILBY default priors π_{B} , with weights expressed as

$$\mathcal{W} = \frac{\pi_{\text{B}}}{\pi_{\text{LI}}}. \quad (\text{C1})$$

We must also account for the fact that BILBY_PIPE uses default priors that are flat in $\mathcal{M}$ and q , whereas LALINFERENCE uses priors that are uniform in component masses. We therefore rejection sample from the released posterior samples with weights given by the inverse of the Jacobian given in equation (21) of Veitch et al. (2015),

$$\mathcal{J} = \frac{\mathcal{M}}{m_1^2}. \quad (\text{C2})$$

The complete reweighting procedure can be written

$$p_{\pi_{\text{B}}} = \mathcal{W}\mathcal{J}p_{\pi_{\text{LI}}}, \quad (\text{C3})$$

where $p_{\pi_{\text{B}}}$ and $p_{\pi_{\text{LI}}}$ are the posterior probabilities computed using BILBY and LALINFERENCE priors, respectively. In practice, we reweight by rejection sampling in order to preserve the independence of samples. We also account for a difference in the definition of the Solar mass $M_{\odot}$ between the current version of BILBY and the version of LALINFERENCE used to produce the public GWTC-1 samples that we compare against.

APPENDIX D: CDF COMPARISONS FOR GWTC-1 EVENTS

In Figures D1–D6, we present the comparisons of the CDFs obtained using BILBY and LALINFERENCE for all parameters and for all events. The legend shows the JS divergence and uncertainty for each parameter, and the shaded regions represent the 1σ , 2σ , and 3σ confidence intervals.

Figure D1. CDF comparison between BILBY and LALINFERENCE for GW150914 and GW151012.

Figure D2. CDF comparison between BILBY and LALINFERENCE for GW151226 and GW170104.

Figure D3. CDF comparison between BILBY and LALINFERENCE for GW170608 and GW170729.

Figure D4. CDF comparison between BILBY and LALINFERENCE for GW170809 and GW170814.

Figure D5. CDF comparison between BILBY and LALINFERENCE for GW170817 and GW170818.

Figure D6. CDF comparison between BILBY and LALINFERENCE for GW170823.**APPENDIX E: PARAMETER DEFINITIONS**

BILBY is able to sample in a range of different parametrizations of compact binaries. In Table E1, we describe the definitions of these parameters as implemented in BILBY. Unless otherwise specified all of these parameters can be sampled in, using the standard waveform model, likelihood, and conversion functions.

Currently, there is a relative lack of support for sampling parameters describing eccentric orbits: the eccentricity e and the argument of periapsis ω . This is because the frequency-domain eccentric waveforms available in LALSIMULATION are less complete than their quasi-circular counterparts, containing only the inspiral section of the signal.

Table E1. Definition of parameters typically considered for CBC inference. Subscript $i = 1, 2$ indicates whether the parameter pertains to the primary (1) or secondary (2) binary object. Subscript $k = x, y, z$ refers to a quantity measured in the $\hat{x}$, $\hat{y}$, or $\hat{z}$ direction; $\hat{z}$ points along the binary axis of rotation, while the $\hat{x}$, $\hat{y}$ directions are orthogonal to each other and $\hat{z}$, defined at reference phase ϕ , and differ by phase offset ϕ_{12} between the two objects. Additional subscripts: * – defined at a reference frequency, † – parameter cannot be sampled, only generated in post-processing, × – parameter cannot yet be sampled or generated in post-processing.

Name	Description	LATEX label	Units
mass _i	Detector-frame (redshifted) mass of the i th object	m_i	$M_\odot$
chirp _{mass}	Detector-frame chirp mass $\mathcal{M} = (m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}$ (Finn & Chernoff 1993; Poisson & Will 1995; Blanchet et al. 1995)	$\mathcal{M}$	$M_\odot$
total _{mass}	Detector-frame combined mass of the primary and secondary masses	M	$M_\odot$
mass _{ratio}	The ratio of the secondary and primary masses $q = m_2/m_1 \leq 1$	q	–
symmetric _{mass_ratio}	A definition of mass ratio which is independent of the identity of the primary/secondary $\eta = q/(1+q)^2$	η	–
mass _{i.source}	Source-frame mass of the i th object $m_i^{\text{source}} = m_i/(1+z)$ (Krolak & Schutz 1987)	m_i^{source}	$M_\odot$
chirp _{mass.source}	Source-frame chirp mass $\mathcal{M}^{\text{source}} = \mathcal{M}/(1+z)$	$\mathcal{M}^{\text{source}}$	$M_\odot$
total _{mass.source}	Source-frame total mass $M^{\text{source}} = M/(1+z)$	M^{source}	$M_\odot$
a _i	Dimensionless spin magnitude of the i th object	a_i	–
tilt _i *	Zenith angle between the spin and orbital angular momenta for the i th object	θ_i	rad
cos _{tilt_i} *	Cosine of the zenith angle between the spin and orbital angular momenta for the i th object	$\cos \theta_i$	–
phi _{j1} *	Difference between total and orbital angular momentum azimuthal angles	ϕ_{j1}	rad
phi ₁₂ *	Difference between the azimuthal angles of the individual spin vector projections on to the orbital plane	ϕ_{12}	rad
chi _i *	i th object aligned spin: projection of the i th object spin on to the orbital angular momentum $\chi_i = a_i \cos(\theta_i)$	χ_i	–

Table E1 – continued

Name	Description	LATEX label	Units
chi.i.in.plane [†]	<i>i</i> th object in-plane spin: magnitude of the projection of the <i>i</i> th object spin on to the orbital plane $\chi_i^\perp = a_i \sin(\theta_i) $	$\chi_i^\perp$	–
chi_eff [†]	Effective inspiral spin parameter $\chi_{\text{eff}} = (\chi_1 + q\chi_2)/(1 + q)$ (Santamaría et al. 2010; Ajith et al. 2011)	χ_{eff}	–
chi_p [†]	Effective precession spin parameter $\chi_p = \max\{\chi_1^\perp, q(3q + 4)/(4q + 3)\chi_2^\perp\}$ (Hannam et al. 2014; Schmidt, Ohme & Hannam 2015)	χ_p	–
spin.i.k [†]	<i>k</i> th component of <i>i</i> th object spin in Euclidean coordinates	$S_{i,k}$	–
lambda.i	Dimensionless tidal deformability of the <i>i</i> th object	Λ_i	–
lambda.tilde	Combined dimensionless tidal deformability (Flanagan & Hinderer 2008; Favata 2014)	$\tilde{\Lambda}$	–
delta.lambda.tilde	Relative difference in the combined tidal deformability (Favata 2014; Wade et al. 2014)	$\delta\tilde{\Lambda}$	–
eccentricity ^{*(†)}	Orbital eccentricity defined at a reference frequency	e	–
argument.of.periaapsis*	The angle between the secondary mass and the ascending node of the orbit when the secondary mass is at periaapsis	ω	rad
x	Right ascension	α	rad
dec	Declination	δ	rad
zenith	Zenith angle in the detector-based sky parametrization	κ	rad
azimuth	Azimuthal angle in the detector-based sky parametrization	ϵ	rad
luminosity.distance	Luminosity distance to the source	d_L	Mpc
comoving.distance	Comoving distance depending on specified cosmology	d_C	Mpc
redshift	Redshift depending on specified cosmology	z	–
geocent.time	GPS reference time at the geocenter, typically merger time	t_c	s
IFO.time	GPS reference time at the detector with name IFO, e.g. H1.time, typically merger time	t_{IFO}	s
time.jitter	Shift to apply for time array used in time marginalization	δt	s
psi	Polarization angle of the source	ψ	rad
phase*	Binary phase at a reference frequency	ϕ	rad
theta.jn	Zenith angle between the total angular momentum and the line of sight	θ_{JN}	rad
cos.theta.jn	Cosine of the zenith angle between the total angular momentum and the line of sight	$\cos \theta_{\text{JN}}$	–
iota*	Zenith angle between the orbital angular momentum and the line of sight	ι	rad
cos.iota*	Cosine of the zenith angle between the orbital angular momentum and the line of sight	$\cos \iota$	–

¹School of Physics and Astronomy, Monash University, Clayton, VIC 3800, Australia

²OzGrav: The ARC Centre of Excellence for Gravitational Wave Discovery, Clayton, VIC 3800, Australia

³LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, USA

⁴LIGO Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁵Department of Physics and Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

⁶School of Physics and Astronomy, Gravity Exploration Institute, Cardiff University, Cardiff CF24 3AA, UK

⁷Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Northwestern University, 1800 Sherman Avenue, Evanston, IL 60201, USA

⁸SUPA, University of Glasgow, Glasgow G12 8QQ, UK

⁹Nikhef, Science Park 105, NL-1098 XG Amsterdam, the Netherlands

¹⁰Institute for High-Energy Physics, University of Amsterdam, Science Park 904, NL-1098 XH Amsterdam, the Netherlands

¹¹Department of Physics, Indian Institute of Technology, Kanpur 208016, India

¹²Inter-University Centre for Astronomy and Astrophysics, Pune 411007, India

¹³Department of Physics, Lancaster University, Lancaster LA1 4YB, UK

¹⁴Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn, VIC 3122, Australia

¹⁵OzGrav: The ARC Centre of Excellence for Gravitational Wave Discovery, Hawthorn, VIC 3122, Australia

¹⁶CNRS, Astroparticule et Cosmologie, Université de Paris, F-75006 Paris, France

¹⁷School of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455, USA

¹⁸Department of Physics, Indian Institute of Technology Bombay, Mumbai, Maharashtra 400076, India

¹⁹Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA

²⁰School of Physics, Georgia Institute of Technology, Atlanta, GA 30332, USA

²¹IAC3-IEEC, Universitat de les Illes Balears, E-07122 Palma de Mallorca, Spain

²²Max Planck Institute for Gravitational Physics (Albert Einstein Institute), D-14476 Potsdam-Golm, Germany

²³Department of Physics, University of Oregon, Eugene, OR 97403, USA

²⁴University of Wisconsin-Milwaukee, Milwaukee, WI 53201, USA

²⁵Institut de Física d'Altes Energies (IFAE), Barcelona Institute of Science and Technology, and ICREA, E-08193 Barcelona, Spain

²⁶CSIRO Astronomy and Space Science, Australia Telescope National Facility, Epping, NSW 1710, Australia

²⁷Department of Physics, University of Florida, Gainesville, FL 32611, USA

²⁸Department of Physics, Università di Roma La Sapienza, I-00185 Roma, Italy

²⁹INFN, Sezione di Roma, I-00185 Roma, Italy

³⁰School of Physics, University of Melbourne, Parkville, VIC 3010, Australia

³¹OzGrav: The ARC Centre of Excellence for Gravitational Wave Discovery, University of Melbourne, Parkville, VIC 3010, Australia

³²Institute for Cosmic Ray Research, The University of Tokyo, 5-1-5 Kashiwanoha, Kashiwa, Chiba 277-8582, Japan

³³National Institute for Mathematical Sciences, Daejeon 34047, South Korea

³⁴School of Physics and Astronomy and Institute for Gravitational Wave Astronomy, University of Birmingham, Edgbaston, Birmingham B15 9TT, UK

This paper has been typeset from a T_EX/L^AT_EX file prepared by the author.