

PyCBC Inference: A Python-based parameter estimation toolkit for compact binary coalescence signals

C. M. Biwer

Los Alamos National Laboratory, Los Alamos, NM 87545, USA
Department of Physics, Syracuse University, Syracuse, NY 13244, USA

Collin D. Capano

Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-30167
Hannover, Germany

Soumi De

Department of Physics, Syracuse University, Syracuse, NY 13244, USA

Miriam Cabero

Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-30167
Hannover, Germany

Duncan A. Brown

Department of Physics, Syracuse University, Syracuse, NY 13244, USA

Alexander H. Nitz

Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-30167
Hannover, Germany

V. Raymond

Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-14476
Potsdam, Germany
School of Physics and Astronomy, Cardiff University, Cardiff, CF243AA, Wales,
UK

Abstract. We introduce new modules in the open-source PyCBC gravitational-wave astronomy toolkit that implement Bayesian inference for compact-object binary mergers. We review the Bayesian inference methods implemented and describe the structure of the modules. We demonstrate that the PyCBC Inference modules produce unbiased estimates of the parameters of a simulated population of binary black hole mergers. We show that the posterior parameter distributions obtained using our new code agree well with the published estimates for binary black holes in the first LIGO-Virgo observing run.

1. Introduction

The observations of six binary black hole mergers [1, 2, 3, 4], and the binary neutron star merger GW170817 [5] by Advanced LIGO [6] and Virgo [7] have established the field of gravitational-wave astronomy. Understanding the origin, evolution, and physics of gravitational-wave sources requires accurately measuring the properties of detected events. In practice, this is performed using Bayesian inference [8, 9]. Bayesian inference allows us to determine the signal model that is best supported by observations and to obtain posterior probability densities for a model’s parameters, hence inferring the properties of the source. In this paper, we present *PyCBC Inference*; a set of Python modules that implement Bayesian inference in the PyCBC open-source toolkit for gravitational-wave astronomy [10]. PyCBC Inference has been used to perform Bayesian inference for several astrophysical problems, including: testing the black hole area increase law [11]; combining multi-messenger observations of GW170817 to constrain the viewing angle of the binary [12]; and determining that the gravitational-wave observations of GW170817 favor a model where both compact objects have the same equation of state, and measuring the tidal deformabilities and radii of the neutron stars [13].

We provide a comprehensive description of the methods and code implemented in PyCBC Inference. We then demonstrate that PyCBC Inference can produce unbiased estimates of the parameters of a simulated population of binary black holes. We show that PyCBC Inference can recover posterior probability distributions that are in good agreement with the published measurements of the binary black holes detected in the first LIGO-Virgo observing run [1]. This paper is organized as follows: Sec. 2 gives an overview of the Bayesian inference methods used in gravitational-wave astronomy for compact-object binary mergers. We provide an overview of the waveform models used; the likelihood function for a known signal in stationary, Gaussian noise; the sampling methods used to estimate the posterior probability densities and the evidence; the selection of independent samples; and the estimation of parameter values from posterior probabilities. Sec. 3 describes the design of the PyCBC Inference software and how the methods described in Sec. 2 are implemented in the code. Sec. 4 uses a simulated population

of binary black holes and the black-hole mergers detected in the first LIGO-Virgo observing run to demonstrate the use of PyCBC Inference. We provide the posterior probability densities for the events GW150914, GW151226, and LVT151012, and the command lines and configurations to reproduce these results as supplemental materials [14]. Finally, we summarize the status of the code and possible future developments in Sec. 5.

2. Bayesian Inference for Binary Mergers

In gravitational-wave astronomy, Bayesian methods are used to infer the properties of detected astrophysical sources [15, 16, 17, 18]. Given the observed data $\vec{d}(t)$ —here this is data from a gravitational-wave detector network in which a search has identified a signal [19, 20, 21]—Bayes’ theorem [8, 9] states that for a hypothesis H ,

$$p(\vec{\vartheta}|\vec{d}(t), H) = \frac{p(\vec{d}(t)|\vec{\vartheta}, H)p(\vec{\vartheta}|H)}{p(\vec{d}(t)|H)}. \tag{1}$$

In our case, hypothesis H is the model of the gravitational-wave signal and $\vec{\vartheta}$ are the parameters of this model. Together, these describe the properties of the astrophysical source of the gravitational waves. In Eq. (1), the prior probability density $p(\vec{\vartheta}|H)$ describes our knowledge about the parameters before considering the observed data $\vec{d}(t)$, and the likelihood $p(\vec{d}(t)|\vec{\vartheta}, H)$ is the probability of obtaining the observation $\vec{d}(t)$ given the waveform model H with parameters $\vec{\vartheta}$.

Often we are only interested in a subset of the parameters $\vec{\vartheta}$. To obtain a probability distribution on one or a few parameters, we marginalize the posterior probability by integrating $p(\vec{d}(t)|\vec{\vartheta}, H)p(\vec{\vartheta}|H)$ over the unwanted parameters. Marginalizing over all parameters yields the evidence, $p(\vec{d}(t)|H)$, which is the denominator in Eq. (1). The evidence serves as a normalization constant of the posterior probability for the given model H . If we have two competing models H_A and H_B , the evidence can be used to determine which model is favored by the data via the Bayes factor [22, 23, 24],

$$\mathcal{B} = \frac{p(\vec{d}(t)|H_A)}{p(\vec{d}(t)|H_B)}. \tag{2}$$

If \mathcal{B} is greater than 1 then model H_A is favored over H_B , with the magnitude of \mathcal{B} indicating the degree of belief.

PyCBC Inference can compute Bayes factors and produce marginalized posterior probability densities given the data from a network of gravitational-wave observatories with N detectors $\vec{d}(t) = \{d_i(t); 1 < i < N\}$, and a model H that describes the astrophysical source. In the remainder of this section, we review the methods used to compute these quantities.

2.1. Waveform Models

The gravitational waves radiated in a binary merger's source frame are described by the component masses $m_{1,2}$, the three-dimensional spin vectors $\vec{s}_{1,2}$ of the compact objects [25], and the binary's eccentricity e [26]. A parameter ϕ describes the phase of the binary at a fiducial reference time, although this is not usually of physical interest. For binaries containing neutron stars, additional parameters $\Lambda_{1,2}$ describe the star's tidal deformabilities [27, 28], which depend on the equation of state of the neutron stars. The waveform observed by the Earth-based detector network depends on seven additional parameters: the signal's time of arrival t_c , the binary's luminosity distance d_L , and four Euler angles that describe the transformation from the binary's frame to the detector network frame [29]. These angles are typically written as the binary's right ascension α , declination δ , a polarization angle Ψ , and the inclination angle ι (the angle between the binary's angular momentum axis and the line of sight).

Binary mergers present a challenging problem for Bayesian inference, as the dimensionality of the signal parameter space is large. This is further complicated by correlations between the signal's parameters. For example, at leading order the gravitational waveform depends on the chirp mass \mathcal{M} [30]. The mass ratio enters the waveform at higher orders and is more difficult to measure. This results in an amplitude-dependent degeneracy between the component masses [18]. Similarly, the binary's mass ratio can be degenerate with its spin [31], although this degeneracy can be broken if the binary is precessing. Much of the effort of parameter estimation in gravitational-wave astronomy has focused on developing computationally feasible ways to explore this signal space, and on extracting physically interesting parameters (or combinations of parameters) from the large, degenerate parameter space (see e.g. Ref. [32] and references therein). However, in many problems of interest, we are not concerned with the full parameter space described above. For example, field binaries are expected to have negligible eccentricity when they are observed by LIGO and Virgo [30], and so eccentricity is neglected in the waveform models. Simplifying assumptions can be made about the compact object's spins (e.g. the spins are aligned with the binary's orbital

angular momentum), reducing the dimensionality of the waveform parameters space.

Given a set of parameters $\vec{\vartheta}$, one can obtain a model of the gravitational-wave signal from a binary merger using a variety of different methods, including: post-Newtonian theory (see e.g. Ref. [33] and references therein), analytic models calibrated against numerical simulations [34, 35, 36, 37, 38, 39, 40], perturbation theory [41, 42], and full numerical solution of the Einstein equations (see e.g. Ref. [43] and references therein). Obtaining posterior probabilities and evidences can require calculating $\mathcal{O}(10^9)$ template waveforms, which restricts us to models that are computationally efficient to calculate. The cost of full numerical simulations makes them prohibitively expensive at present. Even some analytic models are too costly to be used, and surrogate models have been developed that capture the features of these waveforms at reduced computational cost [44, 45].

The specific choice of the waveform model H for an analysis depends on the physics that we wish to explore, computational cost limitations, and the level of accuracy desired in the model. A variety of waveform models are available for use in PyCBC Inference, either directly implemented in PyCBC or via calls to the LIGO Algorithm Library (LAL) [46]. We refer to the PyCBC and LAL documentation, and references therein, for detailed descriptions of these models. In this paper, we demonstrate the use of PyCBC Inference using the IMRPhenomPv2 [47, 48] waveform model for binary black hole mergers. This model captures the inspiral-merger-ringdown physics of spinning, precessing binaries and parameterizing spin effects using a spin magnitude a_j , an azimuthal angle θ_j^a , and a polar angle θ_j^p for each of the two compact objects. Examples of using PyCBC Inference with different waveform models include the analysis of Ref. [13] that used the TaylorF2 post-Newtonian waveform model with tidal corrections, and Ref. [11] that used a ringdown-only waveform that models the quasi-normal modes of the remnant black hole.

2.2. Likelihood Function

The data observed by the gravitational-wave detector network enters Bayes' theorem through the likelihood $p(\vec{d}(t)|\vec{\vartheta}, H)$ in Eq. (1). Currently, PyCBC Inference assumes that the each detector produces stationary, Gaussian noise $n_i(t)$ that is uncorrelated between the detectors in the network. The observed data is then $d_i(t) = n_i(t) + s_i(t)$, where $s_i(t)$ is the gravitational waveform observed in the i -th detector. For detectors that are *not* identical and co-located (as in the case of the LIGO-Virgo network), each detector observes a slightly different waveform due to their different antennae patterns, however the signal in the i -th

detector can be calculated given the subset of the parameters $\vec{\vartheta}$ that describes the location of the binary.

Under these assumptions, the appropriate form of $p(\vec{d}(t)|\vec{\vartheta}, H)$ is the well-known likelihood for a signal of known morphology in Gaussian noise (see e.g. Ref. [49] for its derivation), which is given by

$$p(\vec{d}(t)|\vec{\vartheta}, H) = \exp \left[-\frac{1}{2} \sum_{i=1}^N \langle \tilde{n}_i(f) | \tilde{n}_i(f) \rangle \right] \\ = \exp \left[-\frac{1}{2} \sum_{i=1}^N \langle \tilde{d}_i(f) - \tilde{s}_i(f, \vec{\vartheta}) | \tilde{d}_i(f) - \tilde{s}_i(f, \vec{\vartheta}) \rangle \right], \quad (3)$$

where N is the number of detectors in the network. The inner product $\langle \tilde{a} | \tilde{b} \rangle$ is

$$\langle \tilde{a}_i(f) | \tilde{b}_i(f) \rangle = 4\Re \int_0^\infty \frac{\tilde{a}_i(f) \tilde{b}_i^*(f)}{S_n^{(i)}(f)} df, \quad (4)$$

where $S_n^{(i)}(f)$ is the power spectral density of the of the i -th detector's noise. Here, $\tilde{d}_i(f)$ and $\tilde{n}_i(f)$ are the frequency-domain representations of the data and noise, obtained by a Fourier transformation of $d_i(t)$ and $n_i(t)$, respectively. The model waveform $\tilde{s}_i(f, \vec{\vartheta})$ may be computed directly in the frequency domain, or in the time domain and then Fourier transformed to the frequency domain. There are several operations (e.g. Fourier transforms, noise power spectral density estimation, and inner products) that are common between the calculation of Eq. (3) and the computation of the matched filter signal-to-noise ratio (SNR) in PyCBC [19, 20, 21]. PyCBC Inference uses these existing functions, where appropriate.

In general, gravitational-wave signals consist of a superposition of harmonic modes. However, in many cases it is sufficient to model only the most dominant mode, since the sub-dominant harmonics are too weak to be measured. In this case, the signal observed in all detectors has the same simple dependence on the fiducial phase ϕ ,

$$\tilde{s}_i(f, \vec{\vartheta}, \phi) = \tilde{s}_i^0(f, \vec{\vartheta}, 0) e^{i\phi}. \quad (5)$$

The posterior probability $p(\vec{\vartheta}|\vec{d}(t), H)$ can be analytically marginalized over ϕ for such models [49]. Assuming a uniform prior on $\phi \in [0, 2\pi)$, the marginalized posterior is

$$\log p(\vec{\vartheta}|\vec{d}(t), H) \propto \log p(\vec{\vartheta}|H) + I_0 \left(\left| \sum_i O(\tilde{s}_i^0, \tilde{d}_i) \right| \right) \\ - \frac{1}{2} \sum_i \left[\langle \tilde{s}_i^0, \tilde{s}_i^0 \rangle - \langle \tilde{d}_i, \tilde{d}_i \rangle \right], \quad (6)$$

where

$$\tilde{s}_i^0 \equiv \tilde{s}_i(f, \vec{\vartheta}, \phi = 0),$$

$$O(\tilde{s}_i^0, \tilde{d}_i) \equiv 4 \int_0^\infty \frac{\tilde{s}_i^*(f; \vartheta, 0) \tilde{d}_i(f)}{S_n^{(i)}(f)} df,$$

and I_0 is the modified Bessel function of the first kind.

We have found that analytically marginalizing over ϕ in this manner reduces the computational cost of the analysis by a factor of 2 – 3. The IMRPhenomPv2 model that we use here is a simplified model of precession that allows for this analytic marginalization [47, 48]. Since fiducial phase is generally a nuisance parameter, we use this form of the likelihood function in Secs. 4.1 and 4.2.

2.3. Sampling Methods

Stochastic sampling techniques, and in particular Markov-chain Monte Carlo (MCMC) methods [50, 51, 52, 53], have been used to numerically sample the posterior probability density function of astrophysical parameters for binary-merger signals [18, 54, 55, 56]. Ensemble MCMC algorithms use multiple Markov chains to sample the parameter space. A simple choice to initialize the k -th Markov chain in the ensemble is to draw a set of parameters $\vec{\vartheta}_1^{(k)}$ from the prior probability density function. The Markov chains move around the parameter space according to the following set of rules. At iteration l , the k -th Markov chain has the set of parameters $\vec{\vartheta}_l^{(k)}$. The sampling algorithm chooses a new proposed set of parameters $\vec{\vartheta}_l'^{(k)}$ with probability $Q(\vec{\vartheta}_l'^{(k)}, \vec{\vartheta}_l^{(k)})$. When a new set of parameters is proposed, the sampler computes an acceptance probability γ which determines if the Markov chain should move to the proposed parameter set $\vec{\vartheta}_l'^{(k)}$ such that $\vec{\vartheta}_{l+1}^{(k)} = \vec{\vartheta}_l'^{(k)}$. If $\vec{\vartheta}_l'^{(k)}$ is rejected, then $\vec{\vartheta}_{l+1}^{(k)} = \vec{\vartheta}_l^{(k)}$. After a sufficient number of iterations, the ensemble converges to a distribution that is proportional to a sampling of the posterior probability density function. The true astrophysical parameters $\vec{\vartheta}$ can then be estimated from histograms of the position of the Markov chains in the parameter space. Different ensemble sampling algorithms make particular choices for the proposal probability $Q(\vec{\vartheta}_l'^{(k)}, \vec{\vartheta}_l^{(k)})$ and acceptance probability γ .

The open-source community has several well-developed software packages that implement algorithms for sampling the posterior probability density function. PyCBC Inference leverages these developments, and we have designed a flexible framework that allows the user to choose from multiple ensemble sampling algorithms. Currently, PyCBC Inference supports the open-source ensemble sampler `emcee` [57, 58], its parallel-tempered version `emcee.pt` [59, 60], and the `kombine` [61, 62] sampler. All of the three are ensemble MCMC samplers. The sampling algorithm advances the positions of the walkers based on their previous positions and provides PyCBC Inference the positions of the walkers along the Markov chain.

The `emcee.pt` sampler is a parallel-tempered

sampler which advances multiple ensembles based on the tempering or the “temperatures” used to explore the posterior probability density function. The posterior probability density function for a particular temperature T is modified such that

$$p_T(\vec{\vartheta}|\vec{d}(t), H) = \frac{p(\vec{d}(t)|\vec{\vartheta}, H)^{\frac{1}{T}} p(\vec{\vartheta}|H)}{p(\vec{d}(t)|H)}. \quad (7)$$

The `emcee_pt` sampler uses several temperatures in parallel, and the position of Markov chains are swapped between temperatures using an acceptance criteria described in Ref. [60]. Mixing of Markov chains from the different temperatures makes parallel-tempered samplers suitable for sampling posterior probability density functions with widely separated modes in the parameter space [60]. The `emcee` sampler performs the sampling using one temperature where $T = 1$.

The `kombine` sampler on the other hand uses clustered kernel-density estimates to construct its proposal distribution, and proposals are accepted using the Metropolis–Hastings condition [63]. The `kombine` sampler has been included in PyCBC Inference due to its efficient sampling which significantly lowers the computational cost of an analysis relative to the `emcee_pt` sampler. However, in Sec. 4.1, we found that the nominal configuration of the `kombine` sampler produced biased estimates of parameters for binary black holes.

2.4. Selection of Independent Samples

The output returned by the sampling algorithms discussed in Sec. 2.3 are Markov chains. Successive states of these chains are not independent, as Markov processes depend on the previous state [64]. The autocorrelation length τ_K of a Markov chain is a measure of the number of iterations required to produce independent samples of the posterior probability density function [65]. The autocorrelation length of the k -th Markov chain $X_l^{(k)} = \{\vartheta_g^{(k)}; 1 < g < l\}$ of length l obtained from the sampling algorithm is defined as

$$\tau_K = 1 + 2 \sum_{i=1}^K \hat{R}_i, \quad (8)$$

where K is the first iteration along the Markov chain the condition $m\tau_K \leq K$ is true, m being a parameter which in PyCBC Inference is set to 5 [65]. The autocorrelation function \hat{R}_i is defined as

$$\hat{R}_i = \frac{1}{l\sigma^2} \sum_{t=1}^{l-i} (X_t - \mu)(X_{t+i} - \mu), \quad (9)$$

where X_t are the samples of $X_l^{(k)}$ between the 0-th and the t -th iteration, X_{t+i} are the samples of $X_l^{(k)}$

between the 0-th and the $(t+1)$ -th iterations. Here, μ and σ^2 are the mean and variance of X_t , respectively.

The initial positions of the Markov chains influence their subsequent positions. The length of the Markov chains before they are considered to have lost any memory of the initial positions is called the “burn-in” period. It is a common practice in MCMC analyses to discard samples from the burn-in period to prevent any bias introduced by the initial positions of the Markov chains on the estimates of the parameters from the MCMC. PyCBC Inference has several methods to determine when the Markov chains are past the burn-in period. Here, we describe two methods, `max_posterior` and `n_acl`, which we have found to work well with the `kombine` and `emcee_pt` samplers used in Sections 4.1 and 4.2.

The `max_posterior` algorithm is an implementation of the burn-in test used for the MCMC sampler in Ref. [32]. In this method, the k -th Markov chain is considered to be past the burn-in period at the first iteration l for which

$$\log \mathcal{L}_l^{(k)} \geq \max_{k,l} \log \mathcal{L} - \frac{N_p}{2}, \quad (10)$$

where \mathcal{L} is the prior-weighted likelihood

$$\mathcal{L} = p(\vec{d}(t)|\vec{\vartheta}, H)p(\vec{\vartheta}|H), \quad (11)$$

and N_p is the number of dimensions in the parameter space. The maximization $\max_{k,l} \log \mathcal{L}$ is carried out over all Markov chains and iterations. The ensemble is considered to be past the burn-in period at the first iteration where all chains pass this test. We have found this test works well with the `kombine` sampler if the network signal-to-noise ratio of the signal is $\gtrsim 5$.

While the `max_posterior` test works well with the `kombine` sampler, we have found that it underestimates the burn-in period when used with the `emcee_pt` sampler. Instead we use the `n_acl` test with the `emcee_pt` sampler. This test posits that the sampler is past the burn-in period if the length of the chains exceed 10 times the autocorrelation length. The autocorrelation length is calculated using samples from the second half of the Markov chains. If the test is satisfied, the sampler is considered to be past the burn-in period at the midway point of the Markov chains.

Correlations between the neighboring samples after the burn-in period are removed by “thinning” or drawing samples from the Markov chains with an interval of the autocorrelation length [64]. This is done so that the samples used to estimate the posterior probability density function are independent. Therefore, the number of independent samples of the posterior probability density function is equal to the number of Markov chains used in the ensemble times the number of iterations after the burn-in period divided by the autocorrelation length. PyCBC

Inference will run until it has obtained the desired number of independent samples after the burn-in period.

2.5. Credible Intervals

After discarding samples from the burn-in period and thinning the remaining samples of the Markov chains, the product is the set of independent samples as described in Sec. 2.4. Typically we summarize the measurement of a given parameter using a credible interval. The $x\%$ credible interval is an interval where the true parameter value lies with a probability of $x\%$. PyCBC Inference provides the capability to calculate credible intervals based on percentile values. In the percentile method, the $x\%$ credible interval of a parameter value is written as A_{-B}^{+C} where A is typically the 50-th percentile (median) of the marginalized histograms. The values $A - B$ and $A + C$ represent the lower and upper boundaries of the x -th percentile respectively.

An alternative method of calculating a credible interval estimate is the Highest Posterior Density (HPD) method. An $x\%$ HPD interval is the shortest interval that contains $x\%$ of the probability. The percentile method explained above imposes a non-zero lower boundary to the interval being measured. This can be perceived as a limitation in cases where the weight of histogram at the ~ 0 -th percentile is not significantly different from the weight at the lower boundary of the credible interval. Intervals constructed using the HPD method may be preferred in such cases. Previous studies have noted that HPD intervals may be useful when the posterior distribution is not symmetric [66]. PyCBC Inference uses HPD to construct confidence contours for two-dimensional marginal distributions, but HPD is not used in the construction of one-dimensional credible intervals for a single parameter. This functionality will be included in a future release of PyCBC Inference.

3. The PyCBC Inference Toolkit

In this section we describe the implementation of PyCBC Inference within the broader PyCBC toolkit. PyCBC provides both modules for developing code and executables for performing specific tasks with these modules. The code is available on the public GitHub repository at <https://github.com/gwastro/pycbc>, with executables located in the directory `bin/inference` and the modules in the directory `pycbc/inference`. PyCBC Inference provides an executable called `pycbc_inference` that is the main engine for performing Bayesian inference with PyCBC. A call graph of `pycbc_inference` is shown in Figure 1. In this section, we review the structure of the

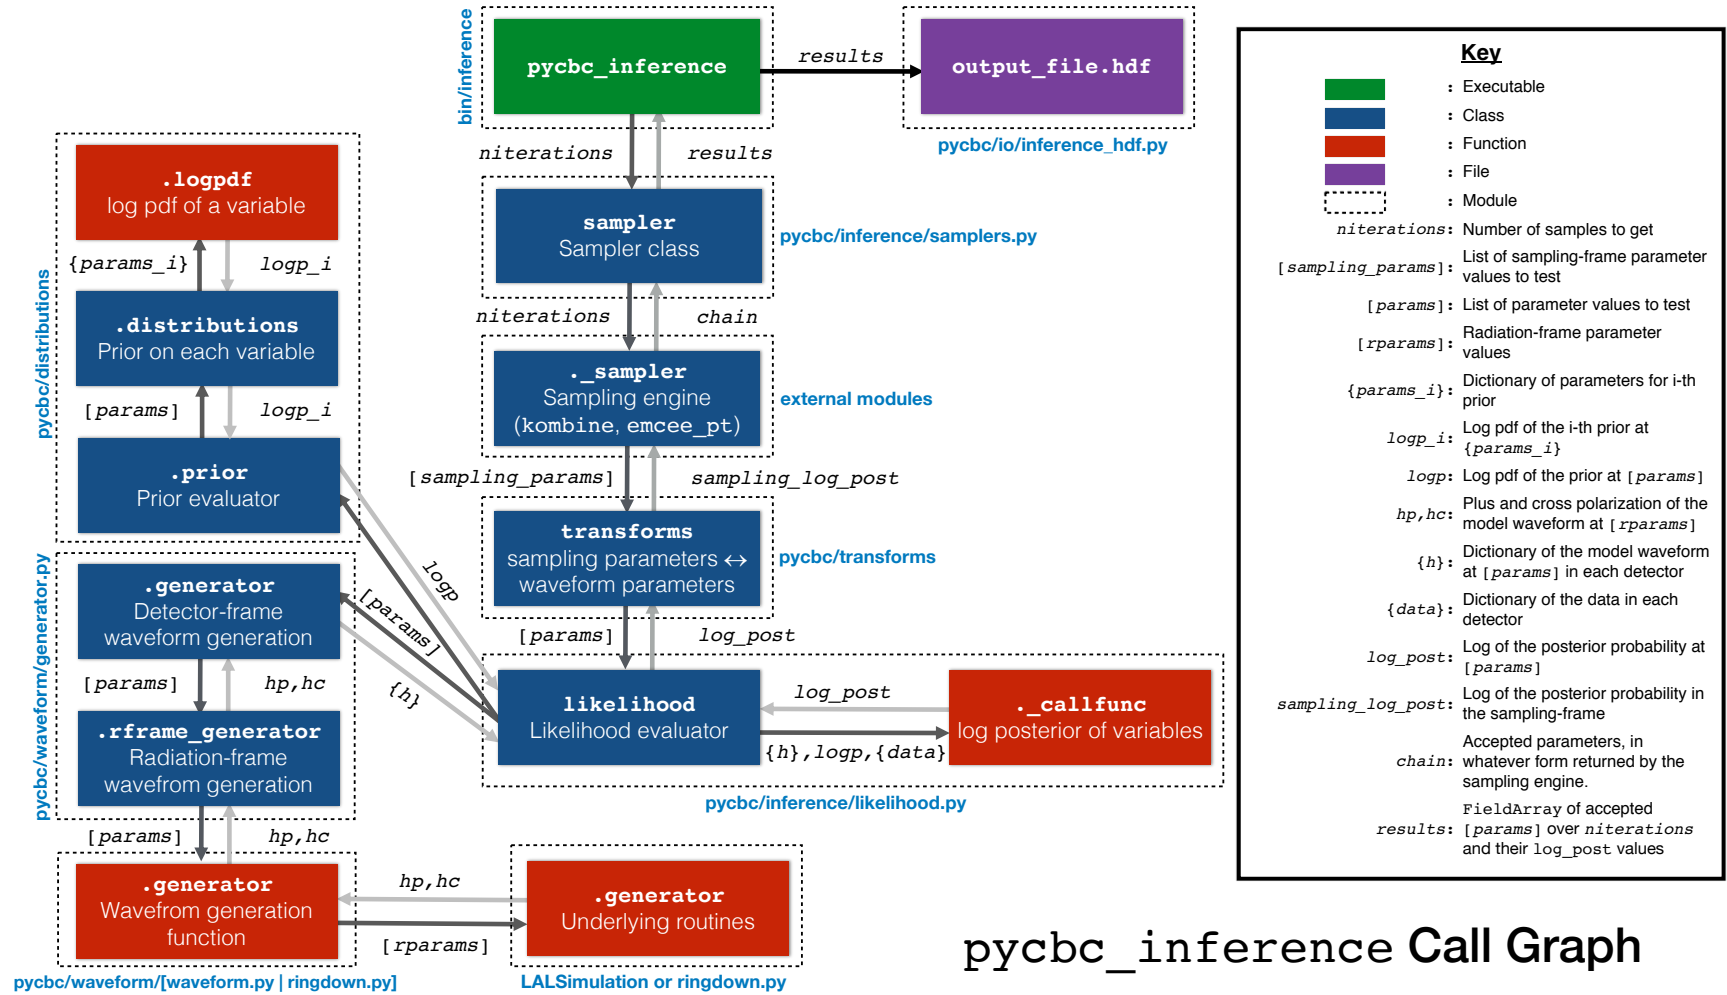
main engine and the Python objects used to build `pycbc_inference`.

3.1. `pycbc_inference` executable

The methods presented in Sec. 2.1, 2.2, 2.3, 2.4, and 2.5 are used to build the executable `pycbc_inference`. For faster performances, `pycbc_inference` can be run on high-throughput computing frameworks such as HTCondor [67, 68] and the processes for running the sampler can be parallelized over multiple compute nodes using MPI [69, 70, 71]. The execution of the likelihood computation and the PSD estimation are done using either single-threaded or parallel FFT engines, such as FFTW [72] or the Intel Math Kernel Library (MKL). For maximum flexibility in heterogeneous computing environments, the processing scheme to be used is specified at runtime as a command line option to `pycbc_inference`.

The input to `pycbc_inference` is a configuration file which contains up to seven types of sections. The `variable_args` section specifies the parameters that are to be varied in the MCMC. There is a `prior` section for each parameter in the `variable_args` section which contains arguments to initialize the prior probability density function for that parameter. There is a `static_args` section specifying any parameter for waveform generation along with its assigned value that should be fixed in the ensemble MCMC. Optionally, the configuration file may also include a `constraint` section(s) containing any conditions that constrain the prior probability density functions of the parameters. For efficient convergence of a Markov chain, it may be desirable to sample the prior probability density function in a different coordinate system than the parameters defined in the `variable_args` section or the parameters inputted to the waveform generation functions. Therefore, the configuration file may contain a `sampling_parameters` and `sampling_transform` section(s) that specifies the transformations between parameters in the `variable_args` sections and the parameters evaluated in the prior probability density function. Finally, the waveform generation functions recognize only a specific set of input parameters. The `waveform_transforms` section(s) may be provided which maps parameters in the `variable_args` section to parameters understood by the waveform generation functions. More details on application of constraints and execution of coordinate transformations are provided in Sec. 3.3 and 3.4 respectively.

The location of the configuration file, gravitational-wave detector data files, data conditioning settings, and settings for the ensemble MCMC are supplied on the command line interface to `pycbc_inference`. The results from running `pycbc_inference` are stored in a



pycbc_inference Call Graph

Figure (1) The executable `pycbc_inference` samples the posterior probability density function. For an iteration in an ensemble MCMC algorithm, the `Sampler` object uses the `LikelihoodEvaluator` object to compute the natural logarithm of the posterior probability, and returns it to `pycbc_inference`. The `LikelihoodEvaluator` object uses the `Generator` object to generate the waveform and `Distribution` objects to evaluate the prior probability density function. Samples are periodically written to the output file.

HDF [73] file whose location is provided on the command line to `pycbc_inference` as well. The main results of interest are stored under the HDF groups `['samples']` and `['likelihood_stats']`. The `['samples']` group contains the history of the Markov chains as separate datasets for each of the variable parameters. The `['likelihood_stats']` group contains a dataset of the natural logarithm of the Jacobian which is needed to transform from the variable parameters to sampling parameters, a dataset containing natural logarithm of the likelihood ratio $\log p(\vec{d}(t)|\vec{\vartheta}, H)/p(\vec{d}(t)|\vec{\eta})$ and a dataset containing the natural logarithm of the prior probabilities. The natural logarithm of the noise likelihood $\log p(\vec{d}(t)|\vec{\eta})$ is stored as an attribute in the output file, and the likelihood is the summation of this quantity with the natural logarithm of the likelihood ratio. Each of the datasets under the `['samples']` group and the `['likelihood_stats']` group has shape `nwalkers × niterations` if the sampling algorithm used in the analysis did not include parallel tempering, and has shape `nsteps × nwalkers × niterations` for parallel-tempered samplers. Here, `nwalkers` is the number of Markov chains, `niterations` is the number of iterations, and `nsteps` is the number of temperatures.

`pycbc_inference` has checkpointing implemented which allows users to resume an analysis from the last set of Markov chains positions written to the output file. It is computationally expensive to obtain the desired number of independent samples using ensemble MCMC methods, and the `pycbc_inference` processes may terminate early due to problems on distributed-computing networks. Therefore, the samples from the Markov chains should be written at regular intervals so `pycbc_inference` can resume the ensemble MCMC from the position of the Markov chains near the state the process was terminated. The frequency `pycbc_inference` writes the samples from Markov chains and the state of the random number generator to the output file and a backup file is specified by the user on the command line. A backup file is written by `pycbc_inference` because the output file from `pycbc_inference` may be corrupted. For example, if the process is aborted while writing to the output file, then the output file may be corrupted. In that case, samples and the state of the random number generator are loaded from the backup file, and the backup file is copied to the output file. This ensures that the `pycbc_inference` process can always be resumed.

For analyses that use the `emcee_pt` sampler, the likelihood can be used to compute the natural logarithm of the evidence using the `emcee_pt` sampler’s `thermodynamic_integration_log_evidence` function [59]. Then, the evidences from two analyses can be used to compute the Bayes factor \mathcal{B} for the comparison

of two waveform models.

We provide example configuration files and run scripts for the analysis of the binary black hole mergers detected in the Advanced LIGO’s first observing run in Ref. [14]. These examples can be used with the open-source datasets provided by the LIGO Open Science Center [74]. The results of these analyses are presented in Sec. 4.2.

3.2. Sampler objects

The PyCBC Inference modules provide a set of `Sampler` objects which execute the Bayesian sampling methods. These objects provide classes and functions for using open-source samplers such as `emcee` [58], `emcee_pt` [59] or `kombine` [62]. This acts as an interface between PyCBC Inference and the external sampler package. The executable `pycbc_inference` initializes, executes, and saves the output from the `Sampler` objects. A particular `Sampler` object is chosen on the command line of `pycbc_inference` with the `--sampler` option. The `Sampler` object provides the external sampler package the positions of the walkers in the parameter space, the natural logarithm of the posterior probabilities at the current iteration, the current “state” determined from a random number generator, and the number of iterations that the sampler is requested to run starting from the current iteration. After running for the given number of iterations, the sampler returns the updated positions of the Markov chains, the natural logarithm of the posterior probabilities, and the new state.

3.3. Transform objects

The `Transform` objects in PyCBC Inference are used to perform transformations between different coordinate systems. Currently, the `Transform` objects are used in two cases: sampling transforms and waveform transforms.

Sampling transforms are used for transforming parameters that are varied in the ensemble MCMC to a different coordinate system before evaluating the prior probability density function. Since there exists degeneracies between several parameters in a waveform model it is useful to parameterize the waveform using a preferred set of parameters which could minimize the correlations. This leads to more efficient sampling, and therefore, it leads to faster convergence of the Markov chains. One example of a sampling transformation is the transformation between the component masses m_1 and m_2 to chirp mass \mathcal{M} and mass ratio q . The convention adopted for q in PyCBC Inference is $q = m_1/m_2$, where m_1 and m_2 are the component masses with $m_1 > m_2$. The chirp mass \mathcal{M} is the most accurately measured parameter in a waveform

model because it is in the leading order term of the post-Newtonian expression of the waveform model. In contrast, the degeneracies of the mass ratio with spin introduces uncertainties in measurements of the component masses. Therefore, sampling in $(\mathcal{M}$ and q) proves to be more efficient than m_1 and m_2 [55, 32, 75]. In the GW150914, LVT151012, and GW151226 configuration files in [14], we demonstrate how to allow the `Sampler` object to provide priors in the (m_1, m_2) coordinates, and specify sampling transformations to the (\mathcal{M}, q) coordinates.

Waveform transforms are used to transform any variable parameters in the ensemble MCMC that may not be understood by the waveform model functions. In PyCBC, the waveform model functions accept the following parameters: component masses m_1 and m_2 , d_L , ι , t_c , ϕ , and any additional spin parameters in Cartesian coordinates. The convention adopted in PyCBC for ι denotes $\iota = 0$ as a “face-on” binary (line of sight parallel to binary angular momentum), $\iota = \frac{\pi}{2}$ as an “edge-on” binary (line of sight perpendicular to binary angular momentum), and $\iota = \pi$ as a “face-off” binary (line of sight anti-parallel to binary angular momentum).

3.4. LikelihoodEvaluator object

The `LikelihoodEvaluator` object computes the natural logarithm of the prior-weighted likelihood given by the numerator of Eq. 1. Since the evidence is constant for a given waveform model, then the prior-weighted likelihood is proportional to the posterior probability density function, and it can be used in sampling algorithms to compute the acceptance probability γ instead of the full posterior probability density function. The prior-weighted likelihood is computed for each new set of parameters as the `Sampler` objects advance the Markov chains through the parameter space.

3.5. Distribution objects

The `LikelihoodEvaluator` object must compute the prior probability density function $p(\vec{\vartheta}|H)$. There exists several `Distribution` objects that provide functions for evaluating the prior probability density function to use for each parameter, and for drawing random samples from these distributions. Currently, PyCBC Inference provides the following `Distributions`:

- (i) `Arbitrary` : Reads a set of samples stored in a HDF format file and uses Gaussian kernel-density estimation [76] to construct the distribution.
- (ii) `CosAngle` : A cosine distribution.
- (iii) `SinAngle` : A sine distribution.
- (iv) `Gaussian` : A multivariate Gaussian distribution.

- (v) `Uniform` : A multidimensional uniform distribution.
- (vi) `UniformAngle` : A uniform distribution between 0 and 2π .
- (vii) `UniformLog` : A multidimensional distribution that is uniform in its logarithm.
- (viii) `UniformPowerLaw` : A multidimensional distribution that is uniform in a power law.
- (ix) `UniformSky` : A two-dimensional isotropic distribution.
- (x) `UniformSolidAngle` : A two-dimensional distribution that is uniform in solid angle.

Multiple `Distribution` objects are needed to define the prior probability density function for all parameters. The `JointDistribution` object combines the individual prior probability density functions, providing a single interface for the `LikelihoodEvaluator` to evaluate the prior probability density function for all parameters. As the sampling algorithm advances the positions of the Markov chains, the `JointDistribution` computes the product of the prior probability density functions for the proposed new set of points in the parameter space. The `JointDistribution` can apply additional constraints on the prior probability density functions of parameters and it renormalizes the prior probability density function accordingly. If multiple constraints are provided, then the union of all constraints are applied. We demonstrate how to apply a cut on \mathcal{M} and q obtained from the m_1 and m_2 prior probability density functions in the GW151226 configuration file in Ref. [14].

3.6. Generator objects

As part of the likelihood calculation described in Sec. 2.2, a waveform $\tilde{s}_i(f, \vec{\vartheta})$ is generated from a waveform model H and set of parameters $\vec{\vartheta}$. PyCBC Inference provides `Generator` objects that allow waveforms $\tilde{s}_i(f, \vec{\vartheta})$ to be generated for waveform models described in Sec. 2.1 using PyCBC’s interface to LAL [46]. There are also `Generator` objects provided for generating ringdown waveforms as used in Ref. [11]. Given the waveform model provided in the configuration file, `pycbc_inference` will automatically select the associated `Generator` object.

4. Validation of the Toolkit

PyCBC Inference includes tools for visualizing the results of parameter estimation, and several analytic functions that can be used to test the generation of known posterior probabilities. Two common ways to visualize results are a scatter plot matrix of the independent samples of the Markov chains, and a marginalized one-dimensional histograms showing the bounds of

each parameter’s credible interval. Analytic likelihood functions available to validate the code include: the multivariate normal, Rosenbrock, eggbox, and volcano functions. An example showing the visualization of results from the multi-variate Gaussian test is shown in Fig. 2. This figure was generated using the executable `pycbc_inference_plot_posterior` which make extensive use of tools from the open-source packages Matplotlib [77] and SciPy [78].

We can also validate the performance of PyCBC Inference by: (i) determining if the inferred parameters of a population of simulated signals agrees with known the parameters that population, and (ii) comparing PyCBC Inference’s parameter credible intervals astrophysical signals to the published LIGO-Virgo results that used a different inference code. In this section, we first check that the credible intervals match the probability of finding the simulated signal parameters in that interval, that is, that $x\%$ of signals should have parameter values in the $x\%$ credible interval. We then compare the recovered parameters of the binary black hole mergers GW150914, GW151226, and LVT151012 to those published in Ref. [1].

4.1. Simulated Signals

To test the performance of PyCBC Inference, we generate 100 realizations of stationary Gaussian noise colored by power-spectral densities representative of the sensitivity of Advanced LIGO detectors at the time of the detection of GW150914 [74]. To each realization of noise we add a simulated signal whose parameters were drawn from the same prior probability density function used in the analysis of GW150914 [79], with an additional cut placed on distance to avoid having too many injections with low matched-filter SNR. The resulting injections have matched-filter SNRs between 5 and 160, with the majority between ~ 10 and ~ 40 . We then perform a parameter estimation analysis on each signal to obtain credible intervals on all parameters.

We perform this test using both the `emcee_pt` and `kombine` samplers. For the `emcee_pt` sampler we use 200 walkers and 20 temperatures. We run the sampler until we obtain at least 2000 independent samples after the burn-in period as determined using the `n_ac1` burn-in test. For the `kombine` sampler, we use 5000 walkers and the `max_posterior` burn-in test. As a result, we need only to run the `kombine` sampler until the burn-in test is satisfied, at which point we immediately have 5000 independent samples of the posterior probability density function.

Both simulated signals and the waveforms in the likelihood computation are generated using IMRPhenomPv2 [47, 48]. This waveform model has 15 parameters. To reduce computational cost, we

analytically marginalize over the fiducial phase ϕ by using Eq. (6) for the posterior probability, thereby reducing the number of sampled parameters to 14. For each parameter, we count the number of times the simulated parameter falls within the measured credible interval.

Figure 3 summarizes the result of this test using the `emcee_pt` and `kombine` samplers. For each of the parameters we plot the fraction of signals whose true parameter value fall within a credible interval as a function of credible interval (this is referred to as a percentile-percentile plot). We expect the former to equal the latter for all parameters, though some fluctuation is expected due to noise. We see that all parameters follow a 1-to-1 relation, though the results from the `kombine` sampler have greater variance than the `emcee_pt` sampler.

To quantify the deviations seen in Fig. 3, we perform a Kolmogorov–Smirnov (KS) test on each parameter to see whether the percentile-percentile curves match the expected 1-to-1 relation. If the samplers and code are performing as expected, then these p-values should in turn follow a uniform distribution. We therefore perform another KS test on the collection of p-values, obtaining a two-tailed p-value of 0.50 for `emcee_pt` and 0.03 for `kombine`. In other words, if `emcee_pt` provides an unbiased estimate of the parameters, then there is a 50% chance that we would obtain a collection of percentile-percentile curves more extreme than seen in Fig. 3. For the `kombine` sampler, the probability of obtaining a more extreme collection of curves than that seen in Fig. 3 is only 3%.

Based on these results, we conclude that PyCBC Inference does indeed provide unbiased estimates of binary black hole parameters when used with `emcee_pt` with the above settings. The `kombine` sampler does not appear to provide unbiased parameter estimates when used to sample the full parameter space of precessing binary black holes with the settings we have used.

4.2. Astrophysical Events

In this section, we present PyCBC Inference measurements of properties of the binary black hole sources of the two gravitational-wave signals GW150914 and GW151226, and the third gravitational-wave signal LVT151012 consistent with the properties of a binary black hole source from Advanced LIGO’s first observing run [80, 1]. We perform the parameter estimation analysis on the Advanced LIGO data available for these events at the LIGO Open Science Center [74]. We use the `emcee_pt` sampler for these analyses. For computing the likelihood, we analyze the gravitational-wave dataset $\vec{d}(t)$ from the Hanford and Livingston detectors. $\vec{d}(t)$ in our analyses are taken from GPS time intervals 1126259452 to

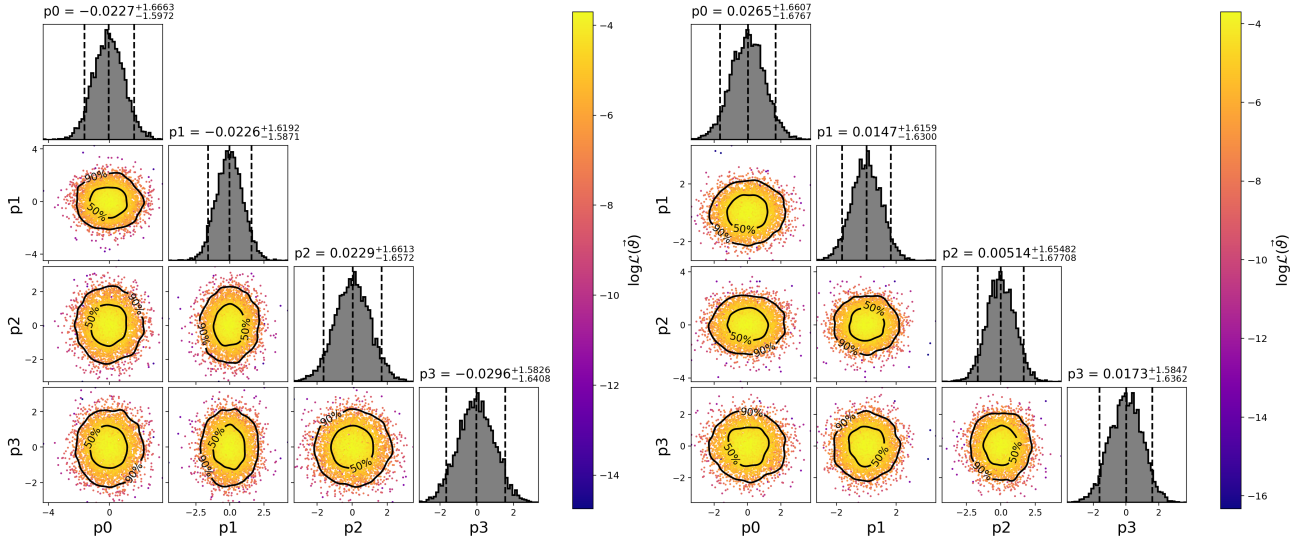


Figure (2) The samples of the posterior probability density function for a four-dimensional normal distribution. Typically, these results are shown as a scatter-plot matrix of independent samples. Here, the points in the scatter-plot matrix are colored by the natural logarithm of the prior-weighted likelihood $\log \mathcal{L}(\vec{\theta})$. At the top of each column is the marginalized one-dimensional histogram for a particular model parameter. In this case, each parameter p_i is the mean of a Gaussian in the range (0, 1). The median and 90% credible interval are superimposed on the marginalized histograms. *Left*: Results obtained from the `emcee_pt` sampler. *Right*: Results obtained from the `combine` sampler.

1126259468 for GW150914, 1135136340 to 1135136356 for GW151226, and 1128678874 to 1128678906 for LVT151012. Detection of gravitational waves from the search pipeline [10, 20, 81, 21, 82] gives initial estimates of the mass, and hence estimates of the length of the signal. From results of the search, LVT151012 was a longer signal with more cycles than the other two events, and LVT151012 had characteristics which were in agreement with a lower mass source than GW150914 and GW151226. Therefore, more data is required for the analysis of LVT151012. The PSD used in the likelihood is constructed using the median PSD estimation method described in Ref. [19] with 8 s Hann-windowed segments (overlapped by 4 s) taken from GPS times 1126258940 to 1126259980 for GW150914, 1135136238 to 1135137278 for GW151226, and 1128678362 to 1128679418 for LVT151012. The PSD estimate is truncated to 4 s in the time-domain using the method described in Ref. [19]. The dataset is sampled at 2048 Hz, and the likelihood is evaluated between a low frequency cutoff of 20 Hz and 1024 Hz.

The waveforms $\tilde{s}_i(f, \vec{\nu})$ used in the likelihood are generated using the IMRPhenomPv2 [47, 48] model implemented in the LIGO Algorithm Library (LAL) [46]. The parameters inferred for these three events are $\vec{\nu} = \{\alpha, \delta, \psi, m_1, m_2, d_L, \iota, t_c, a_1, a_2, \theta_1^a, \theta_2^a, \theta_1^p, \theta_2^p\}$, and we analytically marginalize over the fiducial phase ϕ . These parameters form the complete set of parameters to construct a waveform from a binary black hole

merger, and are the same parameters that were inferred from the parameter estimation analyses in Ref. [1]. Since faster convergence of m_1 and m_2 can be obtained with mass parameterizations of the waveform in \mathcal{M} and q we perform the coordinate transformation from (m_1, m_2) to (\mathcal{M}, q) before evaluating the priors.

We assume uniform prior distributions for the binary component masses $m_{1,2} \in [10, 80] M_\odot$ for GW150914, $m_{1,2} \in [5, 80] M_\odot$ for LVT151012, and $m_{1,2}$ corresponding to chirp mass $\mathcal{M} \in [9.5, 10.5] M_\odot$ and mass ratio $q \in [1, 18]$ for GW151226. We use uniform priors on the spin magnitudes $a_{1,2} \in [0.0, 0.99]$. We use a uniform solid angle prior, where $\theta_{1,2}^a$ is a uniform distribution $\theta_{1,2}^a \in [0, 2\pi)$ and $\theta_{1,2}^p$ is a sine-angle distribution. For the luminosity distance, we use a uniform in volume prior with $d_L \in [10, 1000]$ Mpc for GW150914, $d_L \in [10, 1500]$ Mpc for GW151226, and $d_L \in [10, 2500]$ Mpc for LVT151012. We use uniform priors for the arrival time $t_c \in [t_s - 0.2 s, t_s + 0.2 s]$ where t_s is the trigger time for the particular event obtained from the gravitational-wave search [1, 82]. For the sky location parameters, we use a uniform distribution prior for $\alpha \in [0, 2\pi)$ and a cosine-angle distribution prior for δ . The priors described above are the same as those used in Ref. [1].

The parameter estimation analysis produces distributions that are a sampling of the posterior probability density function for the variable parameters from the ensemble MCMC. We map these distributions

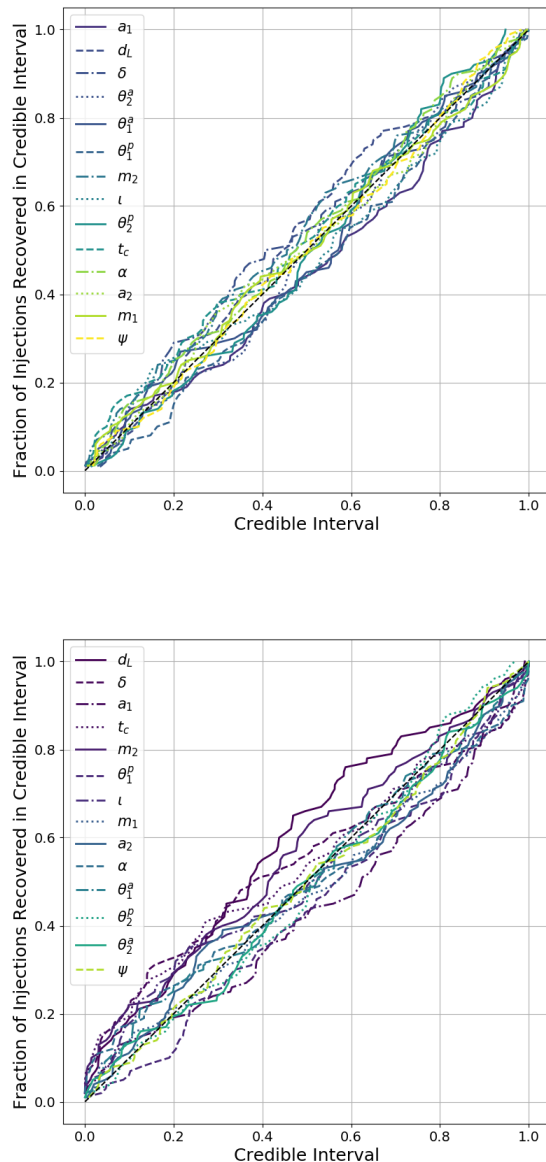


Figure (3) Fraction of simulated signals with parameter values within a credible interval as a function of credible interval. Plotted are all 14 parameters varied in the MCMC analyses. The diagonal line indicates the ideal 1-to-1 relation that is expected if the samplers provide unbiased estimates of the parameters. We perform a Kolmogorov-Smirnov (KS) test on each parameter to obtain two-tailed p-value indicating the consistency between the curves and the diagonal line. *Top*: Results using the `emcee_pt` sampler. *Bottom*: Results using the `kombine` sampler.

obtained directly from the analysis to obtain estimates of other parameters of interest such as the chirp mass \mathcal{M} , mass ratio q , effective spin χ_{eff} , and the precession

spin χ_p [47] parameters. We use d_L to relate the detector-frame masses obtained from the MCMC to the source-frame masses using the standard Λ -CDM cosmology [83, 15].

Recorded in Table 1, is a summary of the median and 90% credible interval values calculated for GW150914, GW151226, and LVT151012 analyses. Results for $m_1^{\text{src}} - m_2^{\text{src}}$, $q - \chi_{\text{eff}}$, and $d_L - \iota$ are shown in Figs. 4, 5, and 6 for GW150914, GW151226, and LVT151012 respectively. The two-dimensional plots in these figures show the 50% and 90% credible regions, and the one-dimensional marginal distributions show the median and 90% credible intervals. Overlaid are the one-dimensional marginal distributions, median, and 90% credible intervals, as well as the 50% and 90% credible regions using the samples obtained from the LIGO Open Science Center [74] for the analyses of the three events reported in Ref. [1] using the IMRPhenomPv2 model for comparison. The results show that GW150914 has the highest mass components among the three events. GW150914 has more support for equal mass ratios whereas the GW151226 and LVT151012 posteriors support more asymmetric mass ratios. Overall, there is preference for smaller spins, with GW151226 having the highest spins among the three events. While the inclination and luminosity distances are not very well constrained, with generally a support for both “face-on” ($\iota = 0$, line of sight parallel to binary angular momentum) and “face-off” ($\iota = \pi$, line of sight anti-parallel to binary angular momentum) systems for all the three events, and GW150914 seem to have more preference for face-off systems. We also computed χ_p for each of the three events and found no significant measurements of precession. Overall, our results are in agreement with those presented in Ref. [1] within the statistical errors of measurement.

5. Conclusions

In this paper we have described PyCBC Inference, a Python-based toolkit with a simplified interface for parameter estimation studies of compact-object binary mergers. We have used this toolkit to estimate the parameters of the gravitational-wave events GW150914, GW151226, and LVT151012; our results are consistent with previously published values. In these analyses, we do not marginalize over calibration uncertainty of the measured strain in our results, which was included in prior work, for example Refs. [1, 79, 84]. We will implement this in PyCBC Inference in the future. We have made the samples of the posterior probability density function from the PyCBC Inference analysis of all three events available in Ref. [14] along with the instructions and configuration files needed

Parameter	GW150914	GW151226	LVT151012
\mathcal{M}^{det}	$31.0^{+1.6}_{-1.5} M_{\odot}$	$9.7^{+0.06}_{-0.06} M_{\odot}$	$18.1^{+1.0}_{-0.7} M_{\odot}$
m_1^{det}	$38.8^{+5.4}_{-3.3} M_{\odot}$	$15.0^{+8.4}_{-3.4} M_{\odot}$	$27.0^{+16.5}_{-5.6} M_{\odot}$
m_2^{det}	$32.9^{+3.2}_{-4.9} M_{\odot}$	$8.4^{+2.3}_{-2.6} M_{\odot}$	$16.3^{+4.2}_{-5.8} M_{\odot}$
\mathcal{M}^{src}	$28.2^{+1.6}_{-1.4} M_{\odot}$	$8.9^{+0.3}_{-0.25} M_{\odot}$	$15.0^{+1.3}_{-1.0} M_{\odot}$
m_1^{src}	$35.3^{+5.0}_{-3.1} M_{\odot}$	$13.7^{+7.7}_{-3.2} M_{\odot}$	$22.4^{+14.1}_{-4.8} M_{\odot}$
m_2^{src}	$29.9^{+3.0}_{-4.4} M_{\odot}$	$7.7^{+2.1}_{-2.4} M_{\odot}$	$13.5^{+3.7}_{-4.7} M_{\odot}$
q	$1.17^{+0.38}_{-0.16}$	$1.78^{+2.21}_{-0.71}$	$1.65^{+2.45}_{-0.6}$
χ_{eff}	$-0.033^{+0.11}_{-0.12}$	$0.2^{+0.18}_{-0.07}$	$0.0023^{+0.24}_{-0.16}$
a_1	$0.29^{+0.57}_{-0.26}$	$0.53^{+0.37}_{-0.45}$	$0.28^{+0.51}_{-0.26}$
a_2	$0.33^{+0.56}_{-0.30}$	$0.51^{+0.43}_{-0.46}$	$0.40^{+0.51}_{-0.36}$
d_L	$497^{+126}_{-202} \text{ Mpc}$	$454^{+164}_{-187} \text{ Mpc}$	$1071^{+458}_{-473} \text{ Mpc}$

Table (1) Results from PyCBC Inference analysis of GW150914, GW151226, and LVT151012. Quoted are the median and 90% credible interval values for the parameters of interest. Interpretations of these results are summarized in Sec. 4.2.

to replicate these results. The source code and documentation for PyCBC Inference is available as part of the PyCBC software package at <http://pycbc.org>.

PyCBC Inference has already been used to produce several astrophysical results: (i) a test of the black hole area increase law [11], (ii) measuring the viewing angle of GW170817 with electromagnetic and gravitational-wave signals [12], and (iii) measuring the tidal deformabilities and radii of neutron stars from the observation of GW170817 [13]. The results presented in this paper and in the studies above demonstrate the capability of PyCBC Inference to perform gravitational-wave parameter estimation analyses. Future developments under consideration are implementation of models to marginalize over calibration errors, generic algorithms to perform model selection, HPD to compute credible intervals, and methods for faster computation of the likelihood.

6. Acknowledgements

The authors would like to thank Will Farr and Ben Farr for valuable insights into the intricacies of ensemble MCMCs. We also thank Ian Harry, Christopher Berry, and Daniel Wysocki for helpful comments on the manuscript. This work was supported by NSF awards PHY-1404395 (DAB, CMB), PHY-1707954 (DAB, SD), and PHY-1607169 (SD). Computations were supported by Syracuse University and NSF award OAC-1541396. We also acknowledge the Max Planck Gesellschaft for support and the Atlas cluster computing team at AEI Hannover. DAB thanks the École de Physique des Houches for hospitality during the completion of this manuscript. The authors thank the LIGO Scientific Collaboration for access to the data and acknowledge the support of the United States National Science Foundation (NSF)

for the construction and operation of the LIGO Laboratory and Advanced LIGO as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, and the Max-Planck-Society (MPS) for support of the construction of Advanced LIGO. Additional support for Advanced LIGO was provided by the Australian Research Council. This research has made use of data obtained from the LIGO Open Science Center <https://losc.ligo.org>.

References

- [1] Abbott B P *et al.* (Virgo, LIGO Scientific) 2016 *Phys. Rev. X* **6** 041015 (*Preprint 1606.04856*)
- [2] Abbott B P *et al.* (VIRGO, LIGO Scientific) 2017 *Phys. Rev. Lett.* **118** 221101 (*Preprint 1706.01812*)
- [3] Abbott B P *et al.* (Virgo, LIGO Scientific) 2017 *Astrophys. J.* **851** L35 (*Preprint 1711.05578*)
- [4] Abbott B P *et al.* (Virgo, LIGO Scientific) 2017 *Phys. Rev. Lett.* **119** 141101 (*Preprint 1709.09660*)
- [5] Abbott B *et al.* (Virgo, LIGO Scientific) 2017 *Phys. Rev. Lett.* **119** 161101 (*Preprint 1710.05832*)
- [6] Aasi J *et al.* (LIGO Scientific) 2015 *Class. Quant. Grav.* **32** 074001 (*Preprint 1411.4547*)
- [7] Acernese F *et al.* (VIRGO) 2015 *Class. Quant. Grav.* **32** 024001 (*Preprint 1408.3978*)
- [8] Bayes M and Price M 1763 *Philosophical Transactions (1683-1775)* URL <http://dx.doi.org/10.1098/rstl.1763.0053>
- [9] Jaynes E T 2003 *Probability Theory: The Logic of Science* (CUP) ISBN 9780521592710 URL <http://www.cambridge.org/au/academic/subjects/physics/theoretical-physics-and-mathematical-physics/probability-theory-logic-science?format=HB&isbn=9780521592710>
- [10] Nitz A *et al.* 2018 *Pycbc v1.9.4* <https://github.com/gwastro/pycbc> URL <https://doi.org/10.5281/zenodo.1208115>
- [11] Cabero M, Capano C D, Fischer-Birnholtz O, Krishnan B, Nielsen A B, Nitz A H and Biwer C M 2018 *Phys. Rev. D* **97**(12) 124069 URL <https://link.aps.org/doi/10.1103/PhysRevD.97.124069>
- [12] Finstad D, De S, Brown D A, Berger E and Biwer C M 2018 *Astrophys. J. Lett.* **860** L2 (*Preprint 1804.04179*)

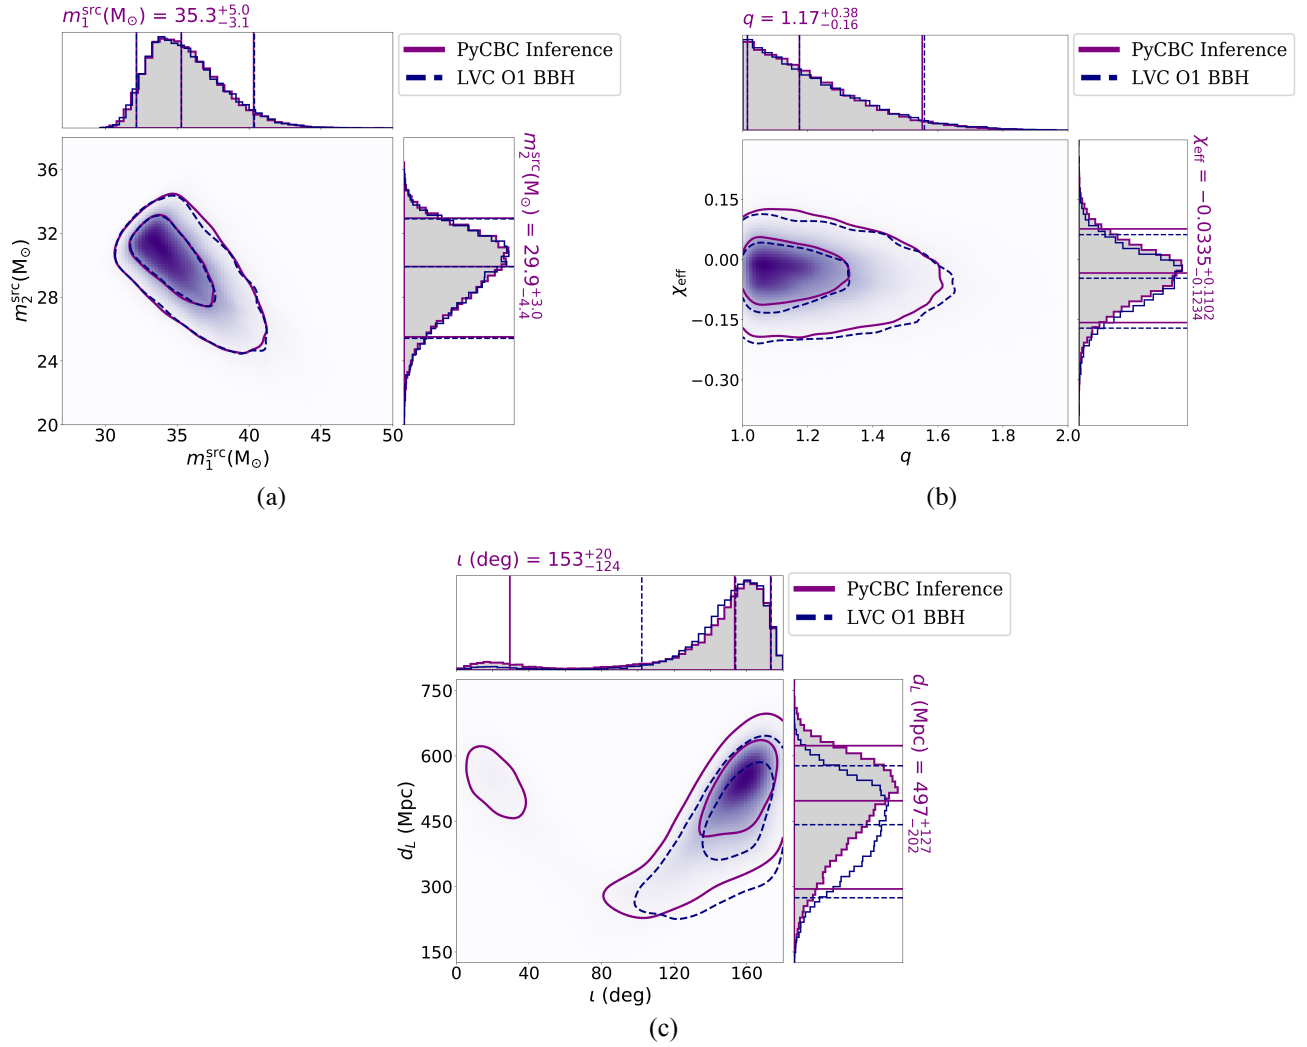


Figure (4) Posterior probability densities for the main parameters of interest from the PyCBC Inference analysis of GW150914. The parameters plotted are (a): $m_1^{\text{src}} - m_2^{\text{src}}$, (b): $q - \chi_{\text{eff}}$ (c): $d_L - i$. The bottom-left panel in each of (a), (b) and (c) show two-dimensional probability densities with 50% and 90% credible contour regions from the PyCBC Inference posteriors. The top-left and the bottom-right panels in each figure show one-dimensional posterior distributions for the individual parameters with solid lines at the 5%, 50% and 95% percentiles. For comparison, we also show 50% and 90% credible regions, one-dimensional posterior probabilities with dashed lines at 5%, 50% and 95% percentiles using the posterior samples obtained from the LIGO Open Science Center [74] for the GW150914 analysis reported in Ref. [1] using the IMRPhenomPv2 model. The measurements show that masses for GW150914 are much better constrained as compared to the other parameters presented. Though there is support for the system being both “face-on” and “face-off”, there seems to be slightly more preference for a ”face-off” system. The posteriors suggest a preference for lower spins. Our measurements are in agreement with the results presented in [1] within the statistical errors of measurement of the parameters.

- [13] De S, Finstad D, Lattimer J M, Brown D A, Berger E and Biwer C M 2018 (*Preprint* 1804.08583)
- [14] <https://github.com/gwastro/pycbc-inference-paper>
- [15] Finn L S and Chernoff D F 1993 *Phys. Rev.* **D47** 2198–2219 (*Preprint* gr-qc/9301003)
- [16] Cutler C and Flanagan E E 1994 *Phys. Rev.* **D49** 2658–2697 (*Preprint* gr-qc/9402014)
- [17] Nicholson D and Vecchio A 1998 *Phys. Rev.* **D57** 4588–4599 (*Preprint* gr-qc/9705064)
- [18] Christensen N and Meyer R 2001 *Phys. Rev.* **D64** 022001 (*Preprint* gr-qc/0102018)
- [19] Allen B, Anderson W G, Brady P R, Brown D A and Creighton J D E 2012 *Phys. Rev.* **D85** 122006
- [20] Usman S A *et al.* 2016 *Class. Quant. Grav.* **33** 215004 (*Preprint* 1508.02357)
- [21] Nitz A H, Dent T, Dal Canton T, Fairhurst S and Brown D A 2017 *Astrophys. J.* **849** 118 (*Preprint* 1705.01513)
- [22] Kass R E and Raftery A E 1995 *Journal of the American Statistical Association* **90** 773–795 (*Preprint* <https://www.tandfonline.com/doi/pdf/10.1080/01621459.1995.10476572>) URL <https://www.tandfonline.com/doi/abs/10.1080/>

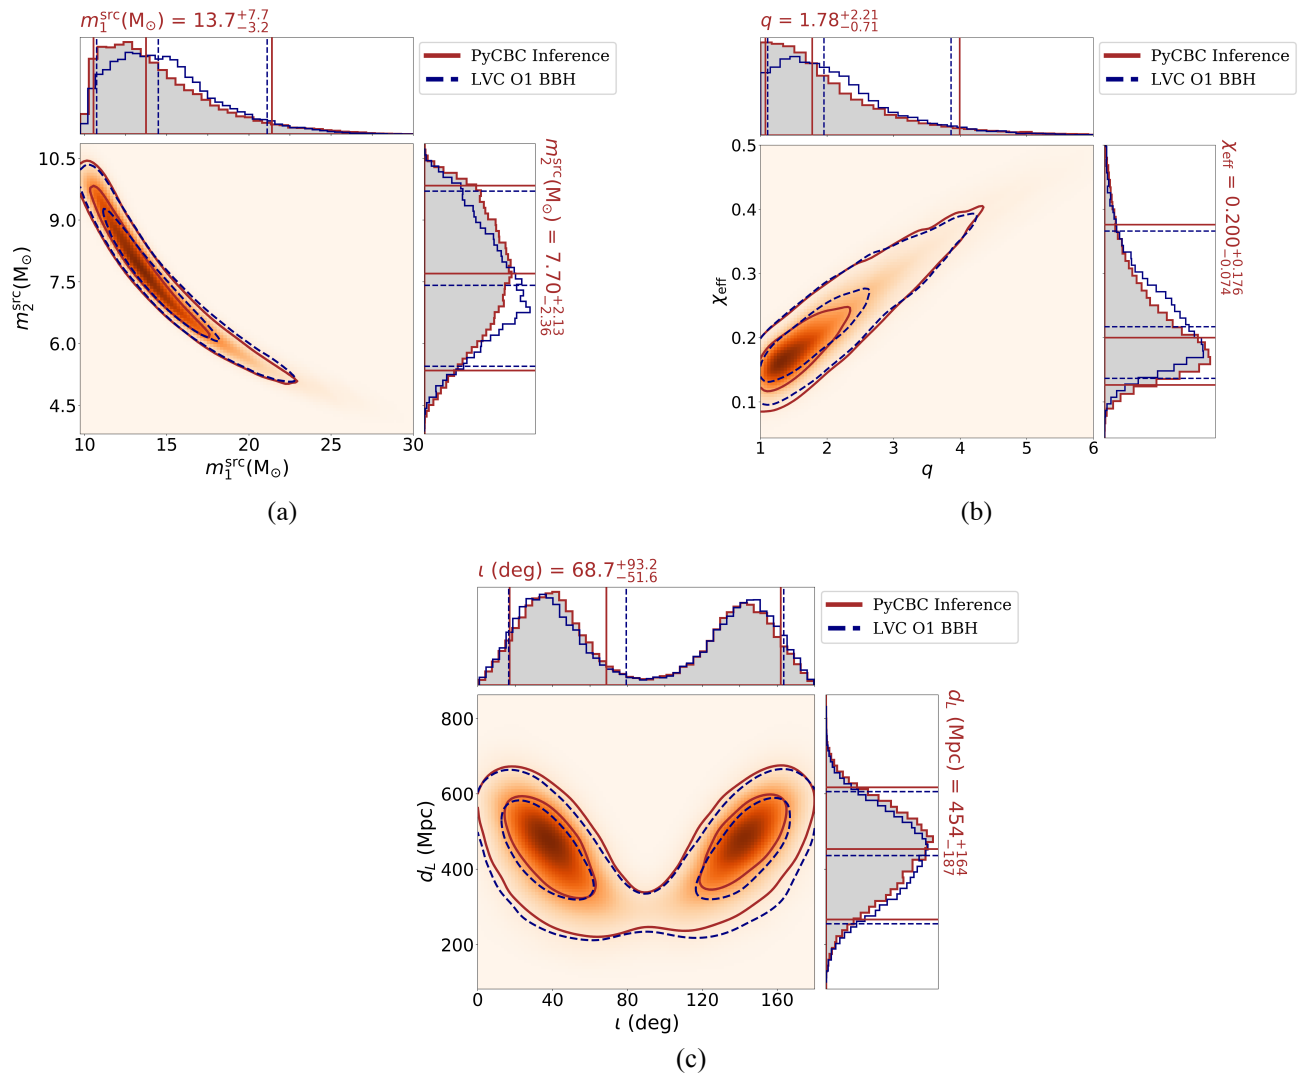


Figure (5) Posterior probability densities for the main parameters of interest from the PyCBC Inference analysis of GW151226. The parameters plotted are (a): $m_1^{\text{src}} - m_2^{\text{src}}$, (b): $q - \chi_{\text{eff}}$ (c): $d_L - \iota$. The bottom-left panel in each of (a), (b) and (c) show two-dimensional probability densities with 50% and 90% credible contour regions from the PyCBC Inference posteriors. The top-left and the bottom-right panels in each figure show one-dimensional posterior distributions for the individual parameters with solid lines at the 5%, 50% and 95% percentiles. For comparison, we also show 50% and 90% credible regions, one-dimensional posterior probabilities with dashed lines at 5%, 50% and 95% percentiles using the posterior samples obtained from the LIGO Open Science Center [74] for the GW151226 analysis reported in Ref. [1] using the IMRPhenomPv2 model. The measurements show that GW151226 is the lowest mass and fastest spinning binary among the three O1 events presented in this work. The posteriors support asymmetric mass ratios. Inclination ι and distance d_L are not well constrained, and there is support for the system being both “face-on” and “face-off”. Our measurements are in agreement with the results presented in [1] within the statistical errors of measurement of the parameters.

- 01621459.1995.10476572
- [23] 1998 *Statistical Science* **13** 163–185 ISSN 08834237 URL <http://www.jstor.org/stable/2676756>
- [24] Skilling J 2006 *Bayesian Anal.* **1** 833–859 URL <https://doi.org/10.1214/06-BA127>
- [25] Thorne K S 1987 THREE HUNDRED YEARS OF GRAVITATION ed Hawking S W and Israel W
- [26] Peters P C 1964 *Phys. Rev.* **136** B1224–B1232
- [27] Flanagan E E and Hinderer T 2008 *Phys. Rev.* **D77** 021502
- [28] Hinderer T 2008 *Astrophys. J.* **677** 1216–1220
- [29] Wahlquist H 1987 *Gen. Rel. Grav.* **19** 1101–1113
- [30] Peters P C and Mathews J 1963 *Phys. Rev.* **131**(1) 435–440 URL <https://link.aps.org/doi/10.1103/PhysRev.131.435>
- [31] Hannam M, Brown D A, Fairhurst S, Fryer C L and Harry I W 2013 *Astrophys. J.* **766** L14 (*Preprint* 1301.5616)
- [32] Veitch J *et al.* 2015 *Phys. Rev.* **D91** 042003 (*Preprint* 1409.7215)
- [33] Blanchet L 2006 *Living Reviews in Relativity* **9** 4 ISSN 1433-8351 URL <https://doi.org/10.12942/lrr-2006-4>

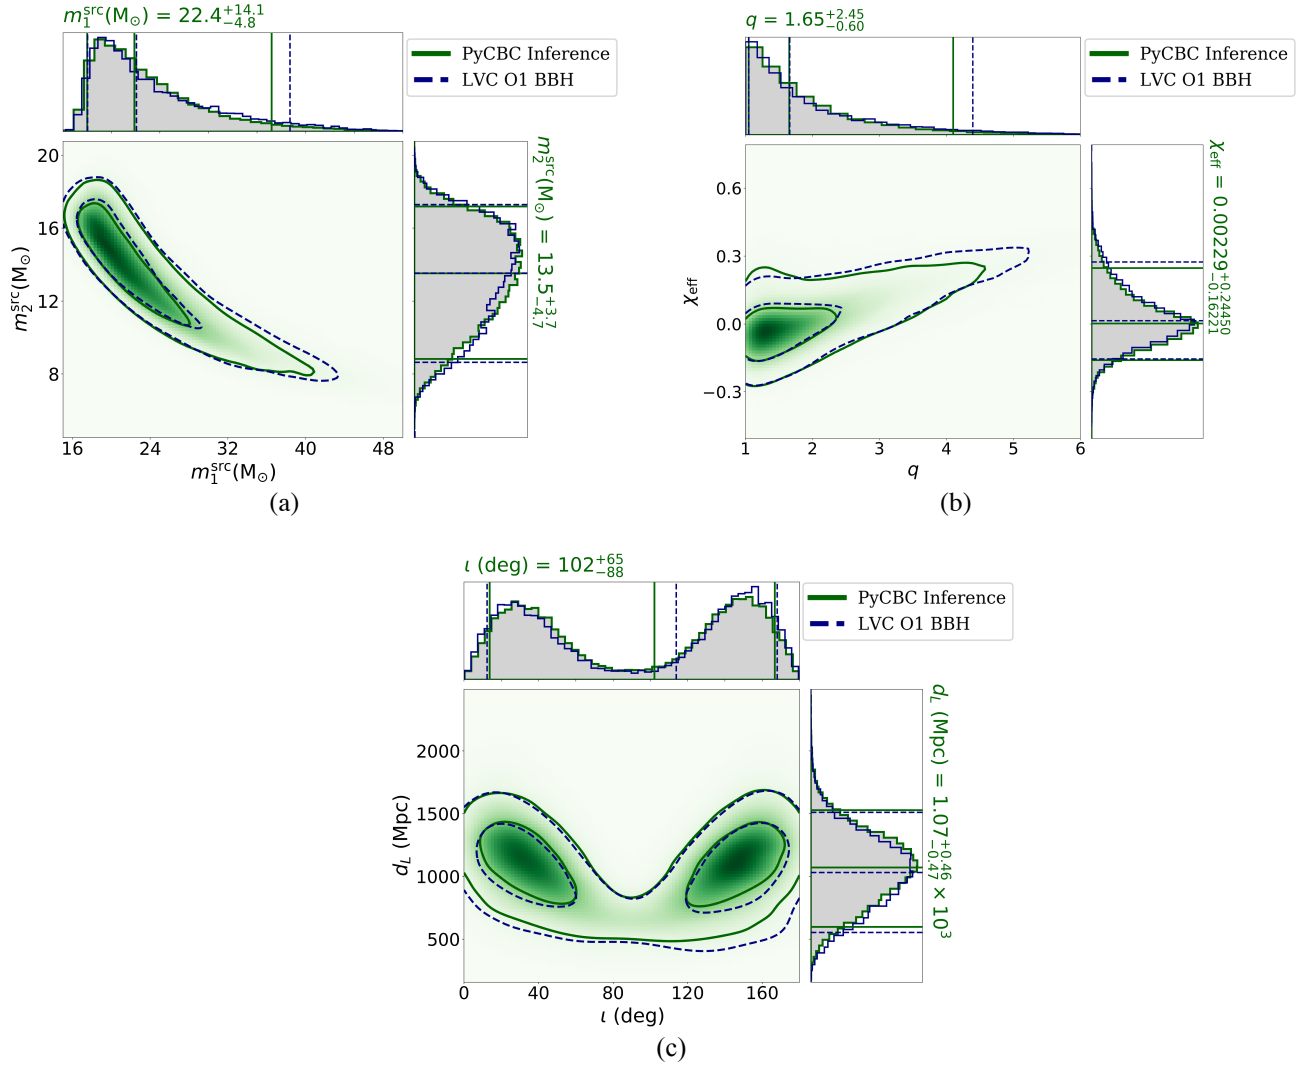


Figure (6) Posterior probability densities for the main parameters of interest from the PyCBC Inference analysis of LVT151012. The parameters plotted are (a): $m_1^{\text{src}} - m_2^{\text{src}}$, (b): $q - \chi_{\text{eff}}$ (c): $d_L - \iota$. The bottom-left panel in each of (a), (b) and (c) show two-dimensional probability densities with 50% and 90% credible contour regions from the PyCBC Inference posteriors. The top-left and the bottom-right panels in each figure show one-dimensional posterior distributions for the individual parameters with solid lines at the 5%, 50% and 95% percentiles. For comparison, we also show 50% and 90% credible regions, one-dimensional posterior probabilities with dashed lines at 5%, 50% and 95% percentiles using the posterior samples obtained from the LIGO Open Science Center [74] for the LVT151012 analysis reported in Ref. [1] using the IMRPhenomPv2 model. The measurements again show that the spins, inclination, and distance are not very well constrained and there is support for the system being both “face-on” and “face-off”. Our measurements are in agreement with the results presented in [1] within the statistical errors of measurement of the parameters.

- [34] Buonanno A and Damour T 1999 *Phys. Rev.* **D59** 084006 (Preprint [gr-qc/9811091](#))
- [35] Buonanno A and Damour T 2000 *Phys. Rev.* **D62** 064015 (Preprint [gr-qc/0001013](#))
- [36] Damour T, Jaranowski P and Schaefer G 2000 *Phys. Rev.* **D62** 084011 (Preprint [gr-qc/0005034](#))
- [37] Damour T 2001 *Phys. Rev.* **D64** 124013 (Preprint [gr-qc/0103018](#))
- [38] Ajith P *et al.* 2007 *Class. Quant. Grav.* **24** S689–S700 (Preprint 0704.3764)
- [39] Ajith P *et al.* 2011 *Phys. Rev. Lett.* **106** 241101 (Preprint 0909.2867)
- [40] Santamaria L *et al.* 2010 *Phys. Rev.* **D82** 064016 (Preprint 1005.3306)
- [41] Teukolsky S A 1972 *Phys. Rev. Lett.* **29** 1114–1118
- [42] Berti E, Cardoso V and Starinets A O 2009 *Class. Quant. Grav.* **26** 163001 (Preprint 0905.2975)
- [43] Cardoso V, Gualtieri L, Herdeiro C and Sperhake U 2015 *Living Rev. Relativity* **18** 1 (Preprint 1409.0014)
- [44] Prer M 2016 *Phys. Rev.* **D93** 064041 (Preprint 1512.02248)
- [45] Lackey B D, Bernuzzi S, Galley C R, Meidam J and Van

- Den Broeck C 2017 *Phys. Rev.* **D95** 104036 (*Preprint* 1610.04742)
- [46] Mercer R A *et al.* 2017 LIGO Algorithm Library <https://git.ligo.org/lscsoft/lalsuite>
- [47] Schmidt P, Ohme F and Hannam M 2015 *Phys. Rev.* **D91** 024043 (*Preprint* 1408.1810)
- [48] Hannam M, Schmidt P, Boh A, Haegel L, Husa S, Ohme F, Pratten G and Prer M 2014 *Phys. Rev. Lett.* **113** 151101 (*Preprint* 1308.3271)
- [49] Wainstein L A and Zubakov V D 1962 *Extraction of signals from noise* (Englewood Cliffs, NJ: Prentice-Hall)
- [50] Metropolis N, Rosenbluth A W, Rosenbluth M N, Teller A H and Teller E 1953 *J. Chem. Phys.* **21** 1087–1092
- [51] Geman S and Geman D 1984 *IEEE Transactions on Pattern Analysis and Machine Intelligence* **PAMI-6** 721–741 ISSN 0162-8828
- [52] Gilks W R, Richardson S and Spiegelhalter D J *Markov Chain Monte Carlo in Practice*, Chapman and Hall, London, 1996
- [53] Gelman A, Robert C, Chopin N and Rousseau J 1995 Bayesian data analysis
- [54] Christensen N, Libson A and Meyer R 2004 *Class. Quant. Grav.* **21** 317–330
- [55] Rover C, Meyer R and Christensen N 2006 *Class. Quant. Grav.* **23** 4895–4906 (*Preprint* gr-qc/0602067)
- [56] Rover C, Meyer R and Christensen N 2007 *Phys. Rev.* **D75** 062004 (*Preprint* gr-qc/0609131)
- [57] Foreman-Mackey D, Hogg D W, Lang D and Goodman J 2013 *Publ. Astron. Soc. Pac.* **125** 306 (*Preprint* 1202.3665)
- [58] Foreman-Mackey D *et al.* 2018 dfm/emcee: emcee v3.0rc1 URL <https://doi.org/10.5281/zenodo.1297477>
- [59] Vousden W *et al.* 2015 URL <https://github.com/willvousden/ptemcee>
- [60] Vousden W D, Farr W M and Mandel I 2016 *Monthly Notices of the Royal Astronomical Society* **455** 1919–1937 (*Preprint* /oup/backfile/content_public/journal/mnras/455/2/10.1093_mnras_stv2422/2/stv2422.pdf) URL <http://dx.doi.org/10.1093/mnras/stv2422>
- [61] Farr B and Farr W M 2015 In prep
- [62] Farr B *et al.* 2018 URL <https://github.com/bfarr/kombine>
- [63] Hastings W K 1970 *Biometrika* **57** 97–109
- [64] Christensen N, Dupuis R J, Woan G and Meyer R 2004 *Phys. Rev.* **D70** 022001 (*Preprint* gr-qc/0402038)
- [65] Madras N and Sokal A D 1988 *J. Statist. Phys.* **50** 109–186
- [66] Chen M H, Shao Q M and Ibrahim J G 2000 *Computing Bayesian Credible and HPD Intervals* (New York, NY: Springer New York) ISBN 978-1-4612-1276-8 URL https://doi.org/10.1007/978-1-4612-1276-8_7
- [67] Tannenbaum T, Wright D, Miller K and Livny M 2001 Condor – a distributed job scheduler *Beowulf Cluster Computing with Linux* ed Sterling T (MIT Press)
- [68] Thain D, Tannenbaum T and Livny M 2005 *Concurrency - Practice and Experience* **17** 323–356
- [69] Dalcin L D, Paz R R, Kler P A and Cosimo A 2011 *Advances in Water Resources* **34** 1124–1139
- [70] Dalcín, Lisandro, Paz R, Storti M and D’Elia J 2008 *Journal of parallel and distributed computing* **68** 655–662
- [71] Dalcín L, Paz R and Storti M 2005 *J. Parallel Distrib. Comput.* **65** 1108–1115 ISSN 0743-7315 URL <http://dx.doi.org/10.1016/j.jpdc.2005.03.010>
- [72] Frigo M and Johnson S G 2005 *Proceedings of the IEEE* **93** 216–231 special issue on “Program Generation, Optimization, and Platform Adaptation”
- [73] Collette A *et al.* 2018 h5py/h5py 2.8.0 URL <https://doi.org/10.5281/zenodo.1246321>
- [74] Vallisneri M, Kanner J, Williams R, Weinstein A and Stephens B 2015 *J. Phys. Conf. Ser.* **610** 012021 (*Preprint* 1410.4839)
- [75] Farr B *et al.* 2016 *Astrophys. J.* **825** 116 (*Preprint* 1508.05336)
- [76] https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.gaussian_kde.html
- [77] Hunter J D 2007 *Computing In Science & Engineering* **9** 90–95
- [78] Jones E, Oliphant T, Peterson P *et al.* 2001– SciPy: Open source scientific tools for Python [Online; accessed jtoday_i] URL <http://www.scipy.org/>
- [79] Abbott B P *et al.* (Virgo, LIGO Scientific) 2016 *Phys. Rev. Lett.* **116** 241102 (*Preprint* 1602.03840)
- [80] Abbott B P *et al.* (Virgo, LIGO Scientific) 2016 *Phys. Rev. Lett.* **116** 061102 (*Preprint* 1602.03837)
- [81] Dal Canton T *et al.* 2014 *Phys. Rev.* **D90** 082004 (*Preprint* 1405.6731)
- [82] Abbott B P *et al.* (Virgo, LIGO Scientific) 2016 *Phys. Rev.* **D93** 122003 (*Preprint* 1602.03839)
- [83] Schutz B F 1986 *Nature* **323** 310–311
- [84] Abbott T D *et al.* (Virgo, LIGO Scientific) 2016 *Phys. Rev.* **X6** 041014 (*Preprint* 1606.01210)