Prompt sky localization of compact binary sources using a meshfree approximation

Lalit Pathak[1,*] Sanket Munishwar[2,†] Amit Reza[3,4,5,‡] and Anand S. Sengupta[1,§]

[1]*Indian Institute of Technology Gandhinagar, Palaj Gandhinagar, Gujarat 382055, India*
[2]*Indian Institute of Science Education and Research Tirupati, Andhra Pradesh 517501, India*
[3]*Nikhef, Science Park 105, 1098 XG Amsterdam, The Netherlands*
[4]*Institute for Gravitational and Subatomic Physics (GRASP), Utrecht University,*
Princetonplein 1, 3584 CC Utrecht, The Netherlands
[5]*Space Research Institute, Austrian Academy of Sciences, Schmiedlstrasse 6, 8042 Graz, Austria*

The number of gravitational wave signals from the merger of compact binary systems detected in the network of advanced LIGO and Virgo detectors is expected to increase considerably in the upcoming science runs. Once a confident detection is made, it is crucial to reconstruct the source's properties rapidly, particularly the sky position and chirp mass, to follow up on these transient sources with telescopes operating at different electromagnetic bands for multimessenger astronomy. In this context, we present a rapid parameter estimation (PE) method aided by mesh-free approximations to accurately reconstruct properties of compact binary sources from data gathered by a network of gravitational wave detectors. This approach builds upon our previous algorithm [L. Pathak *et al.*, Fast likelihood evaluation using meshfree approximations for reconstructing compact binary sources, Phys. Rev. D **108**, 064055 (2023)] to expedite the evaluation of the likelihood function and extend it to enable coherent network PE in a ten-dimensional parameter space, including sky position and polarization angle. Additionally, we propose an optimized interpolation node placement strategy during the start-up stage to enhance the accuracy of the marginalized posterior distributions. With this updated method, we can estimate the properties of binary neutron star sources in approximately 2.4 (2.7) min for the TaylorF2 (IMRPhenomD) signal model by utilizing 64 CPU cores on a shared memory architecture. Furthermore, our approach can be integrated into existing parameter estimation pipelines, providing a valuable tool for the broader scientific community. We also highlight some areas for improvements to this algorithm in the future, which includes overcoming the limitations due to narrow prior bounds.

I. INTRODUCTION

Gravitational wave (GW) signals emitted by compact binary systems made of neutron stars and black holes are expected to take up to a few minutes to sweep through the sensitive band of the advanced LIGO and Virgo detectors, leading up to the epoch of their eventual cataclysmic merger. The merger of two neutron stars/neutron star–black holes (NSBH) [1–5] can be followed by the emission of electromagnetic (EM) radiation at different wavelengths that fade over different timescales ranging from a few minutes to several months, as seen in the fortuitous detection of the binary neutron star (BNS) system GW170817 [6] by the network of advanced LIGO-Virgo detectors [7–11] in the second science run. This discovery resulted in the first

multiwavelength observation of postmerger EM emissions from the merger of two neutron stars. The successful localization and discovery, in this case, were primarily due to the source's fortunate proximity: GW170817 was relatively close (about 40 Mpc), falling well within the sky- and orientation-averaged ranges of the two LIGO detectors, resulting in the loudest signal ever detected by a GW network. The discovery of the EM counterpart of GW170817 improved understanding of the physics of short gamma-ray bursts [12], confirmed the formation of heavy elements after the merger and other areas of new physics, and provided invaluable opportunities to explore the mysteries of the Universe. Therefore, BNS systems are the prime targets for multimessenger astronomy.

In upcoming observing runs, the LIGO-Virgo-KAGRA detectors [7,9–11] are expected to be sensitive above 10 Hz, enabling a more extensive exploration of cosmic volume compared to their current versions. At design sensitivity [13], these detectors are expected to have a BNS range of

---
*lalit.pathak@iitgn.ac.in
†sanket.m@students.iisertirupati.ac.in
‡areza@nikhef.nl
§asengupta@iitgn.ac.in

024053-1

$\sim$330 Mpc (advanced LIGO), $\sim$150–260 Mpc (advanced Virgo) and $\sim$155 Mpc (KAGRA). Consequently, the detection rate is expected to increase to the point where signals from compact binary mergers can be observed daily. It is projected that there could be $10^{+52}_{-10}/1^{+91}_{-1}$ BNS/NSBH events detected in the fourth observing run [13]. To efficiently observe EM counterparts in the future, it will be crucial to prioritize events based on their chirp masses, as suggested by Margalit and Metzger [14]. Furthermore, future GW events could provide opportunities to study various physical properties of binary systems, such as those with eccentricity [15], spin precession [16,17], and more. As such, prompt localization and source reconstruction would be vital for studying EM counterparts of BNS or NSBH binary systems in the future. Moreover, a fast parameter estimation (PE) technique can facilitate the execution of extensive inference investigations [18], like population inference, within a reasonable time frame, which could be impractical when employing conventional brute-force PE methods.

The standard technique for PE uses Bayesian inference, which involves calculating the likelihood of observing the data under the hypothesis that an astrophysical signal (with some model parameters) is embedded in additive Gaussian noise with zero mean and unit variance. This likelihood calculation involves two computationally expensive parts: first, generating the model (template) waveforms [19–21] at sample points proposed by the sampling algorithm and, thereafter, determining the overlap between the data and the waveform. Generating these templates is the most computationally expensive part of the likelihood calculation, especially for low-mass systems like BNSs. This computational burden is amplified by the improved sensitivity of detectors, resulting in a larger number of waveform cycles within the detectors' sensitive band for BNS systems. Additionally, incorporating physical effects into the waveform can further increase the computational cost of generating the waveform and sampling from the expanded parameter space. Wherever possible, factorizing [22] the log-likelihood function into pieces that exclusively depend on the extrinsic and intrinsic parameters, respectively, can lead to a significant reduction in the number of waveform generation and thereby speed up the PE analysis. For example, such factorizations can be made for nonprecessing waveforms, as considered in this work.

In the current operational framework of the LIGO-Virgo-KAGRA (LVK) Collaboration, the primary tool for rapid sky localization is bayestar [23]—a Bayesian, non-Markov chain Monte Carlo (MCMC) sky localization algorithm. This method exhibits remarkable speed, furnishing posterior probability density distributions across sky coordinates within a few tens of seconds following the detection of a gravitational wave signal from the merger of a compact binary source. However, through a comprehensive Bayesian PE analysis encompassing both intrinsic and extrinsic model parameters, Finstad and Brown [24] have established that the accuracy of sky localization can be

significantly augmented, achieving an enhancement of approximately 14 $\deg^2$ over those obtained from the bayestar algorithm alone. Beyond the evident advantages of sky location precision, a full PE analysis also furnishes a highly accurate estimation of the chirp mass for compact binary systems. This additional information plays a pivotal role in making judicious decisions regarding EM follow-up observations, which would be crucial in future observing runs of the network of advanced detectors, further underscoring the merits of developing fast-PE algorithms. Deep-learning-based sky localization tools (CBC-SkyNet [25]) have been recently developed that can obtain "premerger" sky localization areas that are comparable in accuracy to bayestar.

Numerous rapid-PE algorithms have surfaced in the gravitational wave literature, which revolves around two overarching concepts:

(i) Likelihood-based approaches: The first category involves approaches that rapidly evaluate the likelihood function. This set of methods encompasses various methods such as reduced-order models [26–30], heterodyning (or relative binning) [22,31,32], the simple-PE [33] algorithm, and methods based on Gaussian process regression (GPR) [34] like RIFT [35]. Other techniques, such as adaptive frequency resolution-based likelihood evaluation [36] and mass-spin reparametrization-based rapid PE [37], have also been proposed. Improved algorithms allowing for more efficient reduced-order quadrature bases have been recently proposed [38].

(ii) Likelihood-"free" approaches: The second category consists of methods that bypass direct likelihood evaluation and instead learn the posterior distributions using advanced machine-learning (ML) techniques such as deep learning, normalizing flows, and variational inference [39–42].

Hybrid techniques that combine likelihood heterodyning with tools to enhance the convergence of gradient-based MCMC samplers have also been proposed [43]. Other techniques involving score-based diffusion models [44] to learn an empirical noise distribution directly from the detector data have also been proposed.

In this context, our mesh-free approach aligns with the first category of rapid-PE methods since it is designed to swiftly assess the likelihood. Note that the RIFT [35] method also interpolates the likelihood directly using GPR. In comparison, the mesh-free method first represents the likelihood function on an orthonormal basis and then uses radial basis functions to interpolate the coefficients. So, in spite of the apparent similarity, the context in which the "interpolation" technique is applied is completely different in the two methods. The use of GPR (a supervised training method) to "learn" the likelihood function at different points of the parameter space is in the spirit of ML methods, in contrast with the classical approach followed by the mesh-free method.

In our previous work [45], we introduced the mesh-free method as a means of rapidly inferring parameters for compact binary coalescence (CBC) sources for a single detector by considering only six parameters: the two-component masses, two aligned spins, luminosity distance, and coalescence time. In this study, (a) we extend the method to include additional parameters such as sky location, inclination, and polarization, thereby enabling a coherent multidetector PE analysis. (b) Furthermore, we have modified the radial basis function (RBF) mesh-free interpolation node placement scheme, which directly impacts the accuracy of the PE results. In contrast to our previous work, which utilized uniformly distributed nodes over the sample space, we now employ a combination of multivariate Gaussian and uniform distributions for efficient node placement. (c) Our results demonstrate that, when run on 32(64) cores, the mesh-free method can produce accurate marginalized posteriors for the GW170817 event and locate it in the sky within ~3.4(2.4) min of detecting the event. The resulting posteriors obtained from the mesh-free method are statistically indistinguishable from those obtained using the brute-force method implemented in PyCBC [46]. Please note that the brute-force method is not optimized and is used only as a yardstick for accuracy.

The rest of the paper is structured into the following sections: Sec. II introduces the PE basics, where we briefly discuss the Bayesian approach for inferring GW parameters from the data and define the coherent network likelihood. Section III explains the start-up and online stages and discusses the likelihood evaluation procedure using mesh-free interpolation. Section IV A presents the results of a detailed analysis of the mesh-free method for the GW170817 BNS event. In Sec. IV B, we examine the effectiveness of our method by testing it on simulated events covering a wide range of signal-to-noise ratios (SNRs). Finally, we summarize the results in Sec. V and discuss the limitations of the current implementation (such as restricting the sampler over narrow prior bounds). We suggest some ideas for overcoming this limitation in follow-up studies.

II. PARAMETER ESTIMATION

A. Bayesian inference

Given a stretch of GW strain data $\boldsymbol{d}$ from detectors containing a GW signal $\boldsymbol{h}(\vec{\Lambda})$, embedded in additive Gaussian noise $\boldsymbol{n}$, we want to estimate the posterior distribution $p(\vec{\Lambda}|\boldsymbol{d})$ over the source parameters $\vec{\Lambda}$. The posterior distribution, in turn, is related to the likelihood function $\mathcal{L}(\boldsymbol{d}|\vec{\Lambda})$ via the well-known Bayes theorem,

$$p(\vec{\Lambda}|\boldsymbol{d}) = \frac{\mathcal{L}(\boldsymbol{d}|\vec{\Lambda})p(\vec{\Lambda})}{p(\boldsymbol{d})}, \qquad (1)$$

where $p(\vec{\Lambda})$ is the prior distribution over model parameters $\vec{\Lambda} \equiv \{\vec{\lambda}, \vec{\theta}, t_c\}$. Here $\vec{\lambda}$ is a set of intrinsic parameters such as

the component masses $m_{1,2}$ and dimensionless aligned spins $\chi_{1z,2z}$, while $\vec{\theta}$ represents the set of extrinsic parameters such as the source's sky location, i.e., right ascension (α) and declination (δ), the inclination (ι) of the orbital plane of the binary with respect to the line of sight, the polarization angle (ψ), the luminosity distance (d_L) of the source from Earth, and the geocentric epoch of coalescence t_c.

In principle, given a numerical prescription for calculating the likelihood function $\mathcal{L}(\boldsymbol{d}|\vec{\Lambda})$ under a waveform model and assuming prior distributions over the model parameters, we can evaluate the left-hand side of Eq. (1) at any given point in the sample space up to an overall normalization factor. However, with a large number of parameters $\vec{\Lambda}$ (typically ~15 parameters), evaluating $p(\vec{\Lambda}|\boldsymbol{d})$ on a fine grid over the sample space becomes increasingly tedious and eventually intractable with finite computational resources. Therefore, a more intelligent strategy is employed where one uses stochastic sampling techniques to estimate the posterior distribution $p(\vec{\Lambda}|\boldsymbol{d})$. There are a number of schemes to sample the posterior distribution, such as the MCMC scheme [47] and its variants, nested sampling [48] algorithms, etc. In this paper, we use DYNESTY [49,50], an extensively used Python implementation of the nested sampling algorithm for GW data analysis. In this work, we have only focused on speeding up the PE algorithm by quickly evaluating the likelihood function at any point proposed by the sampling algorithm. Admittedly, another aspect of designing a fast-PE algorithm would involve optimizing the sampling algorithm itself—which is not considered in this work. It is easy to see that such improvements to the sampling algorithm can positively impact the performance of many fast-PE methods (including ours). For example, a significant improvement in computation time has been demonstrated by efficiently populating the parameter space as proposed by the VARAHA sampling technique [51].

B. The likelihood function

Let $\boldsymbol{d}^{(i)}$ be the strain data recorded at the ith detector containing an astrophysical GW signal $\tilde{h}^{(i)}(\vec{\Lambda})$. We pause to remark that complex $\tilde{h}^{(i)}(\vec{\Lambda})$ will denote the frequency-domain Fourier transform of the signal $h^{(i)}(\vec{\Lambda})$. Assuming uncorrelated noise among the $N_{\rm d}$ detectors in the network, the coherent log-likelihood [52] is given by

$$\ln\mathcal{L}(\vec{\Lambda}) = \sum_{i=1}^{N_{\rm d}} \langle \boldsymbol{d}^{(i)} | \tilde{h}^{(i)}(\vec{\Lambda}) \rangle$$
$$- \frac{1}{2} \sum_{i=1}^{N_{\rm d}} \left[\|\tilde{h}^{(i)}(\vec{\Lambda})\|^2 + \|\boldsymbol{d}^{(i)}\|^2 \right]. \qquad (2)$$

For a nonprecessing GW signal model, one can use the relation $\tilde{h}_\times \propto -j\tilde{h}_+$ between the two polarization states to express the signal at the ith detector as

$$
\tilde{h}^{(i)}(\vec{\Lambda}) \equiv \tilde{h}(\vec{\Lambda}, t^{(i)}) = \mathcal{A}^{(i)}\tilde{h}_+(\vec{\lambda}, t^{(i)}),
$$
$$
= \mathcal{A}^{(i)}\tilde{h}_+(\vec{\lambda})e^{-j2\pi f t_c}e^{-j2\pi f \Delta t^{(i)}}. \tag{3}
$$

The complex amplitude of the signal $\mathcal{A}^{(i)}$ depends only on the extrinsic parameters $\vec{\theta} \in \vec{\Lambda}$ through the antenna pattern functions, luminosity distance, and the inclination angle,

$$
\mathcal{A}^{(i)} = \frac{1}{d_L}\left[\frac{1+\cos^2\iota}{2}F_+^{(i)}(\alpha,\delta,\psi)-j\cos\iota F_\times^{(i)}(\alpha,\delta,\psi)\right], \tag{4}
$$

where $F_+^{(i)}(\alpha,\delta,\psi)$ and $F_\times^{(i)}(\alpha,\delta,\psi)$ are, respectively, the "plus" and "cross" antenna pattern functions of the ith detector. The antenna pattern functions of a detector describe the angular response of the detector to incoming GW signals. It arises from contracting the position- and geometry-dependent "detector tensor" [53] with the metric perturbations (GW), thereby mapping the latter to a GW strain amplitude $h^{(i)}(\vec{\Lambda})$ recorded at the detector.

As indicated by Eq. (3), the signal acquires an additional phase difference during its projection from the geocentric frame to the detector frame, which corresponds to a time delay denoted by $\Delta t^{(i)}$. This temporal offset originates due to the relative positioning of the ith detector in relation to Earth's center, and it can be expressed explicitly as follows:

$$
\Delta t^{(i)} \equiv t^{(i)} - t_c = \frac{\vec{x}^{(i)} \cdot \hat{N}(\alpha,\delta)}{c}, \tag{5}
$$

where $\vec{x}^{(i)}$ is a vector pointing from Earth's center to the location of the ith detector, $t^{(i)}$ is the time at the ith detector, and $\hat{N}(\alpha,\delta)$ is the direction of the GW propagation [54]. In our analysis, we consider the log-likelihood function marginalized over the coalescence phase [55] parameter, which can be written as follows:

$$
\ln\mathcal{L}(\vec{\Lambda}|\boldsymbol{d}^{(i)})\Big|_{\phi_c} = \ln I_0\left[\left|\sum_{i=1}^{N_d}\langle\boldsymbol{d}^{(i)}|\tilde{h}^{(i)}(\vec{\Lambda})\rangle\right|\right]
$$
$$
-\frac{1}{2}\sum_{i=1}^{N_d}\left[\|\tilde{h}^{(i)}(\vec{\Lambda})\|^2 + \|\boldsymbol{d}^{(i)}\|^2\right], \tag{6}
$$

where $I_0(\cdot)$ is the modified Bessel function of the first kind. By marginalizing over extrinsic parameters, we effectively reduce the dimensionality of the problem, resulting in accelerated likelihood calculations and enhanced sampling convergence. Substituting Eq. (3) in the above equation, we get

$$
\ln\mathcal{L}(\vec{\Lambda}|\boldsymbol{d}^{(i)})\Big|_{\phi_c} = \ln I_0\left[\left|\sum_{i=1}^{N_d}\mathcal{A}^{(i)*}\langle\boldsymbol{d}^{(i)}|\tilde{h}_+(\vec{\lambda},t^{(i)})\rangle\right|\right]
$$
$$
-\frac{1}{2}\sum_{i=1}^{N_d}\left[\left|\mathcal{A}^{(i)}\right|^2\sigma^2(\vec{\lambda})^{(i)} + \|\boldsymbol{d}^{(i)}\|^2\right], \tag{7}
$$

where $\langle\boldsymbol{d}^{(i)}|\tilde{h}_+(\vec{\lambda},t^{(i)})\rangle$ is the complex overlap integral, while $\sigma^2(\vec{\lambda})^{(i)} \equiv \langle\tilde{h}_+(\vec{\lambda},t^{(i)})|\tilde{h}_+(\vec{\lambda},t^{(i)})\rangle$ is the squared norm of the template $\tilde{h}_+(\vec{\lambda})$. $\sigma^2(\vec{\lambda})^{(i)}$ depends on the noise power spectral density (PSD) of the ith detector. The squared norm of the data vector $\|\boldsymbol{d}^{(i)}\|^2$ remains constant during the PE analysis and, therefore, does not impact the overall "shape" of the likelihood. As such, it can be excluded from the subsequent analysis.

III. MESH-FREE LIKELIHOOD INTERPOLATION

As discussed in our previous work [45], the mesh-free method comprises two distinct stages:
  (1) Start-up stage: During this phase, we create RBF interpolants for the pertinent quantities.
  (2) Online stage: In this stage, we swiftly compute the likelihood function at query points proposed by the sampling algorithm.
Next, we will delve into the various stages that constitute mesh-free interpolation.

A. Start-up stage

During the start-up stage, we generate RBF interpolants for the relevant quantities, which enable rapid likelihood calculations in the online stage. The start-up stage can be further divided into the following parts:
  (i) Placement of RBF nodes over the intrinsic parameter space, λ. These discrete nodes are denoted as $\vec{\lambda}^n$: $n = 1, 2, \ldots, N$.
      Note that we set $\Delta t^{(i)} = 0$ for all the detectors while calculating $\vec{z}^{(i)}(\vec{\lambda}^n)$. Additional time delays $\Delta t^{(i)}$ for the signal to reach the ith detector depending on the source's location in the sky can be taken care of during the online stage while sampling the posterior distribution over the sky location parameters, as explained later.
  (ii) Singular value decomposition (SVD) of the complex time-series matrix $\mathcal{Z}^{(i)}$ at the ith detector, where the nth row is defined as $\vec{z}^{(i)}(\vec{\lambda}^n)$.
  (iii) Generation of RBF interpolating functions for the SVD coefficients and the template norm square σ^2 for each detector.
We now describe some of these steps in greater detail.

1. RBF nodes placement

The accurate construction of mesh-free RBF interpolants relies significantly on the strategy of randomly distributing

the RBF nodes, denoted as $\vec{\lambda}^n$, across the sample space. However, the task of finding an optimal set of nodes poses a considerable challenge. Typically, a hypercube in the sample space is selected around a reference point during the sampling of the posterior distribution. This reference point, commonly known as a trigger or best-matched template $\vec{\lambda}^*$ and trigger time (geocentric) t_{trig} are derived from an upstream search pipeline [56,57].

In our previous work, we opted for a uniform distribution of nodes within the hypercube centered around the reference point (injection in the case of a simulated BNS event). This approach works well for events with low SNRs since the corresponding likelihood function exhibited a relatively flatter profile across the parameter space. However, for high SNR events, where the likelihood profile exhibits a sharp peak in certain regions of the parameter space, uniformly distributing nodes within the hypercube may not adequately cover the support of the posterior distribution, particularly in areas where we seek higher accuracy in interpolated likelihood estimates. As a result, we adopted a combination of nodes drawn from both a multivariate Gaussian distribution and a uniform distribution across the sample space within the hypercube. The multivariate Gaussian distribution is defined by a mean vector $\vec{\mu}$ and a covariance matrix $\boldsymbol{\Sigma}$, and its expression is as follows:

$$p(\vec{\lambda}) \propto \exp\left[-(\vec{\lambda}-\vec{\mu})^T \boldsymbol{\Sigma}^{-1}(\vec{\lambda}-\vec{\mu})\right], \qquad (8)$$

where μ can be selected as the trigger point $\vec{\lambda}^*$, and $\boldsymbol{\Sigma}$ can be computed using the inverse of the Fisher matrix Γ evaluated at $\vec{\lambda}^*$. While the distribution of nodes generated using the Fisher matrix may not perfectly match the actual posterior distribution across the sample space, it does exhibit significant overlap with the true distribution, especially in the vicinity of the likelihood peak.

Since the trigger originates from a search pipeline that uses a discrete set of templates over a grid (also known as a template bank) [58,59], we can select a reference point, denoted as $\vec{\lambda}_{\text{ref}}$, in close proximity to the trigger point $(\vec{\lambda}^*)$, which may have a higher network SNR. We employ an optimization strategy [60] aimed at maximizing the network SNR in order to identify this reference point. Subsequently, a portion of the interpolation nodes are generated from a multivariate Gaussian distribution, $\mathcal{N}(\vec{\lambda}_{\text{ref}}, \boldsymbol{\Sigma})$, where $\boldsymbol{\Sigma}$ is computed from the inverse of the Fisher matrix evaluated at $\vec{\lambda}_{\text{ref}}$, to adequately cover the region around the peak. For regions of the sample space where Gaussian nodes are sparse or possibly nonexistent, we spray a uniform distribution of samples across the hypercube. This approach ensures high accuracy in interpolated likelihood estimates throughout the entire sample space.

For a visual representation illustrating the various combinations of Gaussian and uniform nodes, please refer to Fig. 1. In particular, note that choosing a small fraction of random RBF nodes from the multivariate Gaussian distribution leads to more accurate posterior distributions over certain parameters (see Fig. 8 in Appendix A).

2. SVD of $\mathcal{Z}^{(i)}$ and interpolant generation

After obtaining the appropriate RBF nodes $\vec{\lambda}^n$, we can efficiently compute the time series $\vec{z}^{(i)}(\vec{\lambda}^n) \equiv z^{(i)}(\vec{\lambda}^n, t_c)$ using fast-Fourier-transform-based circular correlations, where t_c are the uniformly spaced discrete-time shifts taken from a specified range around a reference coalescence time t_{trig}. Given that the predominant portion of the posterior support originates from the vicinity of the likelihood peak, we opt for $t_c \in [t_{\text{trig}} \pm 0.15 \text{ s}]$ as the interval for sampling t_c.

Similarly, we compute the template norm square $\sigma^2(\vec{\lambda}^n)^{(i)}$ at all the RBF nodes $\vec{\lambda}^n$.
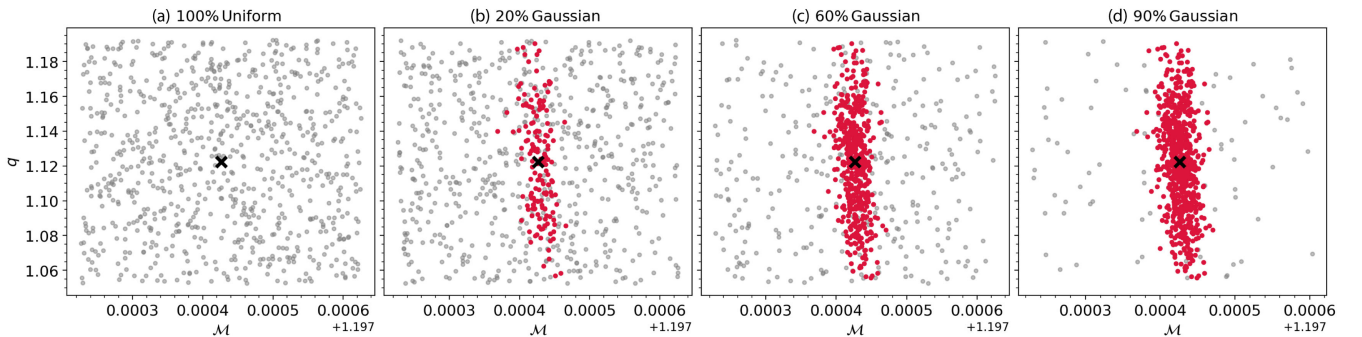

FIG. 1. The figure displays randomly selected RBF interpolation nodes in the intrinsic parameter space with varying fractions from Gaussian and uniform distributions (2D slice in $\mathcal{M}$ and q coordinates). Randomly placed nodes drawn from a Gaussian distribution are denoted in red, while those from a uniform distribution are shown in gray. The covariance matrix of the Gaussian distribution can be estimated semianalytically given the signal model. (a) Represents a choice with 100% (all) nodes drawn from a uniform distribution, and (b)–(d) exhibit mixtures of different percentages of nodes from the uniform and Gaussian distributions. The center (marked as a black **x**) of the sample space is determined through an optimization routine based on the highest network SNR. This search is guided by the best-matched trigger obtained from the upstream search pipeline.

Our objective is to identify the basis vectors capable of spanning the space of $\vec{z}^{(i)}(\vec{\lambda}^n)$. This can be achieved by stacking (row-wise) the time-series vectors $\vec{z}^{(i)}(\vec{\lambda}^n)$ for each $\vec{\lambda}^n$, where $n = 1, 2, ..., N$. Subsequently, performing the SVD on the resultant matrix $\mathcal{Z}^{(i)}$ yields our desired basis vectors $\vec{u}_\mu$. It can be shown that any vector $\vec{z}^{(i)}(\vec{\lambda}^n)$ within the specified parameter range can be expressed as a linear combination of these basis vectors $\vec{u}_\mu$,

$$\vec{z}^{(i)}(\vec{\lambda}^n) = \sum_{\mu=1}^{N} C_\mu^{n(i)} \vec{u}_\mu^{(i)}, \tag{9}$$

where $C_\mu^{n(i)} \equiv C_\mu^{(i)}(\vec{\lambda}^n)$, $\mu = 1, 2, ..., N$ are the N SVD coefficients. These coefficients $C_\mu^{n(i)}$ are characterized as smooth functions over $\vec{\lambda}^n$ within a sufficiently narrow boundary encompassing the posterior support. Therefore, each of these N coefficients can be treated as a scalar-valued smooth function over the d-dimensional intrinsic parameter space and expressed as a linear combination of the RBFs [61] ϕ centered at the interpolation nodes,

$$C_\mu^{q(i)} = \sum_{n=1}^{N} a_n^{(i)} \phi(\|\vec{\lambda}^q - \vec{\lambda}^n\|_2) + \sum_{j=1}^{M} b_j^{(i)} p_j(\vec{\lambda}^q), \tag{10}$$

where ϕ is the RBF kernel centered at $\vec{\lambda}^n \in \mathbb{R}^d$, and $\{p_j\}$ denotes the monomials that span the space of polynomials with a predetermined degree ν in d dimensions. The inclusion of these monomial terms has been shown to enhance the accuracy of RBF interpolation [62].

Similarly, we can also express $\sigma^2(\vec{\lambda}^q)^{(i)}$ in terms of the RBF functions and monomials by treating it as a smoothly varying scalar-valued function over the interpolation domain.

From Eq. (10), we find that there are a total of $(N + M)$ coefficients to be solved for each interpolating function. The SVD coefficients $C_\mu^{q(i)}$ and $\sigma^2(\vec{\lambda}^q)^{(i)}$ are known at each of the N RBF nodes $\vec{\lambda}^n$. These provide N interpolation conditions. To uniquely determine the coefficients a_n and b_j, we impose M additional conditions of the form $\sum_{j=1}^{M} a_j^{(i)} p_j(\vec{\lambda}^q) = 0$. This leads to the following system of $N + M$ equations:

$$\begin{bmatrix} \mathbf{\Phi} & \mathbf{P} \\ \mathbf{P}^T & \mathbf{O} \end{bmatrix} \begin{bmatrix} \mathbf{a}^{(i)} \\ \mathbf{b}^{(i)} \end{bmatrix} = \begin{bmatrix} C_\mu^{n(i)} \\ \mathbf{0} \end{bmatrix}, \tag{11}$$

where the matrices $\mathbf{\Phi}$ and $\mathbf{P}$ have components $\Phi_{ij} = \phi(\|\vec{\lambda}^i - \vec{\lambda}^j\|_2)$ and $P_{ij} = p_j(\vec{\lambda}^i)$, respectively; $\mathbf{O}_{M \times M}$ is a zero matrix and $\mathbf{0}_{M \times 1}$ is a zero vector.

The additional M conditions arise due to the fact that $\mathbf{\Phi}$ needs to be a real, symmetric, and conditionally positive definite matrix of order $\nu + 1$, and the following theorem (reproduced here from [61] for completeness) guarantees the uniqueness of the solutions:

Theorem 1. Let $\mathbf{A}$ be a real symmetric $N \times N$ matrix that is conditionally positive definite of order one, and let $\mathbf{B} = [1, ..., 1]^T$ be an $N \times 1$ matrix (column vector). Then, the system of linear equations

$$\begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^T & \mathbf{O} \end{bmatrix} \begin{bmatrix} \mathbf{c} \\ \mathbf{d} \end{bmatrix} = \begin{bmatrix} \mathbf{y} \\ \mathbf{0} \end{bmatrix}$$

is uniquely solvable.

Although Theorem 1 is applicable for cases where $\nu = 0$, it can be generalized to polynomials of higher degrees (for a proof, refer to Theorem 7.2 in [61]).

Now, the set of equations (11) can be uniquely solved for the unknown coefficients $\mathbf{a}^{(i)}$ and $\mathbf{b}^{(i)}$ to determine the mesh-free interpolating functions corresponding to the coefficients $C_\mu^{(i)}$. From the sharply decreasing spectrum of eigenvalues shown in Fig. 2, it is clear that only the "top-few" basis vectors are sufficient to reconstruct $\vec{z}^{(i)}(\vec{\lambda}^q)$ at minimal reconstruction error; we generate only top-ℓ mesh-free interpolants of $C_\mu^{q(i)}$ where $\mu = 1, ..., \ell$. Note that ℓ is chosen based on the spectrum of eigenvalues of the matrix $\mathcal{Z}^{(i)}$. We have chosen $\ell = 20$ where the normalized eigenvalues fall (see Fig. 2) to $\sim 10^{-12}$. However, this is a subjective choice.

In a similar way, we generate the mesh-free interpolants for the scalar-valued function $\sigma^2(\vec{\lambda}^q)^{(i)}$ as well.

Note that, to optimize computational efficiency, we need to generate a minimum number of RBF nodes $\vec{\lambda}^n$ given the degree (ν) of the polynomial and dimension (d) of the intrinsic parameter space, and it is given by $N_{\min} = \binom{\nu+d}{\nu}$. Increasing the number of nodes increases not only the
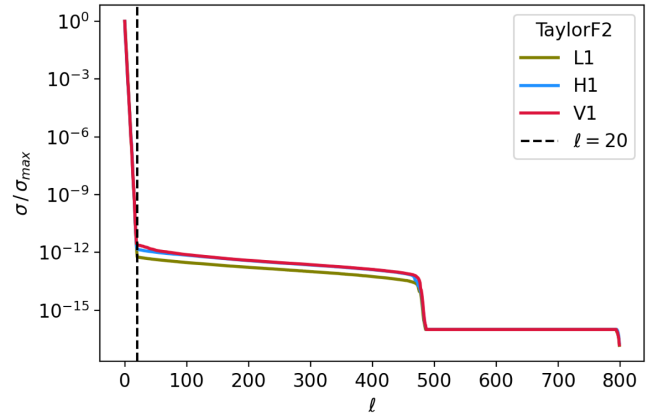

FIG. 2. The figure shows the profile of the singular values normalized by the highest singular value. Notably, there is a sharp decline in the singular values (indicated by a black dashed line) due to the large correlation between the time series $\vec{z}^{(i)}(\vec{\lambda}^n)$ at different RBF interpolation nodes, revealing that only the linear combination of the top-few singular vectors (around 20) is sufficient for effectively reconstructing the time series $\vec{z}^{(i)}(\vec{\lambda}^q)$ with minimal reconstruction error. The 3.5 PN TaylorF2 post-Newtonian signal model is assumed.

computational cost of generating the interpolants but also their evaluation cost during the online stage, which is evident from Eq. (10) as the number of coefficients $a_n^{(i)}$ increase linearly with the number of RBF nodes N. In our study, we empirically choose at least twice (or more) the minimum number of required nodes.

At the end of the start-up stage, we thus have uniquely determined $(\ell + 1)$ interpolating functions.

B. Online stage

Once we have the mesh-free interpolants ready, we can evaluate the mesh-free interpolated values of $C_\mu^{q(i)}$ and $\sigma^2(\vec{\lambda}^q)^{(i)}$ at any query point $\vec{\lambda}^q$ rapidly. We then combine the interpolated values of the coefficients $C_\mu^{q(i)}$ with the corresponding basis vectors $\vec{u}_\mu$ according to Eq. (9), restricting the summation to the first ℓ terms.

In the start-up stage (Sec. III A), we set $\Delta t^{(i)} = 0$ while calculating $\vec{z}^{(i)}(\vec{\lambda}^n)$ and did not include the time delay relative to Earth's center, which depends on the sky location. Now, we incorporate it while reconstructing the interpolated time series $\vec{z}^{(i)}(\vec{\lambda}^q)$. Moreover, we focus on constructing only $\vec{z}^{(i)}(\vec{\lambda}^q)$ at approximately ten time samples centered around query time $t^{q(i)}$ [Eq. (5)]. To ensure subsample accuracy, we fit these samples with a cubic spline, which incurs negligible computational cost. Subsequently, we calculate $\vec{z}^{(i)}(\vec{\lambda}^q)$ at the query $t^{q(i)}$ using the cubic-spline interpolant. Similarly, we evaluate the interpolated value of $\sigma^2(\vec{\lambda}^q)^{(i)}$. Finally, combining $\vec{z}^{(i)}(\vec{\lambda}^q)$ and $\sigma^2(\vec{\lambda}^q)^{(i)}$ with coefficients dependent on extrinsic parameters $\vec{\theta}$, as exemplified in Eq. (7), enables the computation of $\ln \mathcal{L}_{\mathrm{RBF}}$.

IV. NUMERICAL EXPERIMENTS

A. GW170817 event

We conducted a Bayesian PE study on the GW170817 BNS event to test our method. The strain data from the two advanced-LIGO and Virgo detectors were obtained from the open archival datasets available from the Gravitational Wave Open Science Center [63]. We used a 360 s data segment that was cleaned with a fourth-order high-pass filter, with a cutoff frequency of 18 Hz. The seismic cutoff frequency was set to 20 Hz. The TaylorF2 waveform model [64–67] was employed to recover the source parameters. The noise PSD was estimated from the strain data using overlapping segments of length 2 s, using the median-mean PSD estimation method by applying a Hann window as implemented in PyCBC. Subsequently, we generated $N = 800$ RBF nodes using the node placement algorithm described in Sec. III A.

Out of total RBF nodes, $N_{\mathrm{g}} = 200$ nodes (number of Gaussian nodes) are selected from a multivariate Gaussian

distribution $\mathcal{N}(\vec{\lambda}_{\mathrm{ref}}, \boldsymbol{\Sigma})$, where the mean is set as $\vec{\lambda}_{\mathrm{ref}}$, and the covariance matrix $\boldsymbol{\Sigma}$ is estimated by evaluating the inverse of the Fisher matrix $\boldsymbol{\Gamma}$ [68] at the reference point $\vec{\lambda}_{\mathrm{ref}}$ (see Sec. III A 1). We utilize the GWFAST [69] Python package to calculate the covariance matrix. The remaining nodes $N_{\mathrm{u}} = 600$ are uniformly sampled from the ranges specified in Table I. For generating the RBF interpolants, we employ the publicly available RBF Python package [70]. A Gaussian RBF kernel $\phi = \exp(-\varepsilon r^2)$ is used as the basis function, where ε serves as the shape parameter. We set $\varepsilon = 10$ through trial and error, but it can be optimized using the leave-one-out cross-validation [71] method. The degree of the monomials is chosen as $\nu = 7$, and we retain the top 20 basis vectors ($\ell = 20$) for the reconstruction of $\vec{z}^{(i)}(\vec{\lambda}^q)$.

To sample the ten-dimensional parameter space $\vec{\Lambda}$, we utilize nested sampling implemented in the DYNESTY Python package. The likelihood function used for the analysis is $\ln \mathcal{L}_{\mathrm{RBF}}$. The prior distributions and their boundaries for all sampling parameters are provided in Table I. The sampler configuration is outlined as follows: nlive = 500, walks = 100, sample = rwalk, and dlogz = 0.1. These parameters crucially determine the accuracy and time taken by the nested sampling algorithm. Here, nlive represents the number of live points. Opting for a larger number of nlive points results in a more finely sampled posterior (and consequently, evidence), albeit at the expense of requiring more iterations for convergence. Walks denotes the minimum number of points needed before proposing a new live point, sample indicates the chosen sampling approach for generating samples, and dlogz signifies the proportion of the remaining prior volume's contribution to the total evidence, which functions as a stopping criterion for terminating the sampling process. Further details on DYNESTY's nested sampling algorithm and its implementation can be found in [49,50]. For comparison, we also employ the direct marginalized phase likelihood implemented in PyCBC with the DYNESTY sampler, using the same sampler configuration mentioned above.

TABLE I. Prior parameter space over the ten-dimensional parameter space $\vec{\Lambda}$.

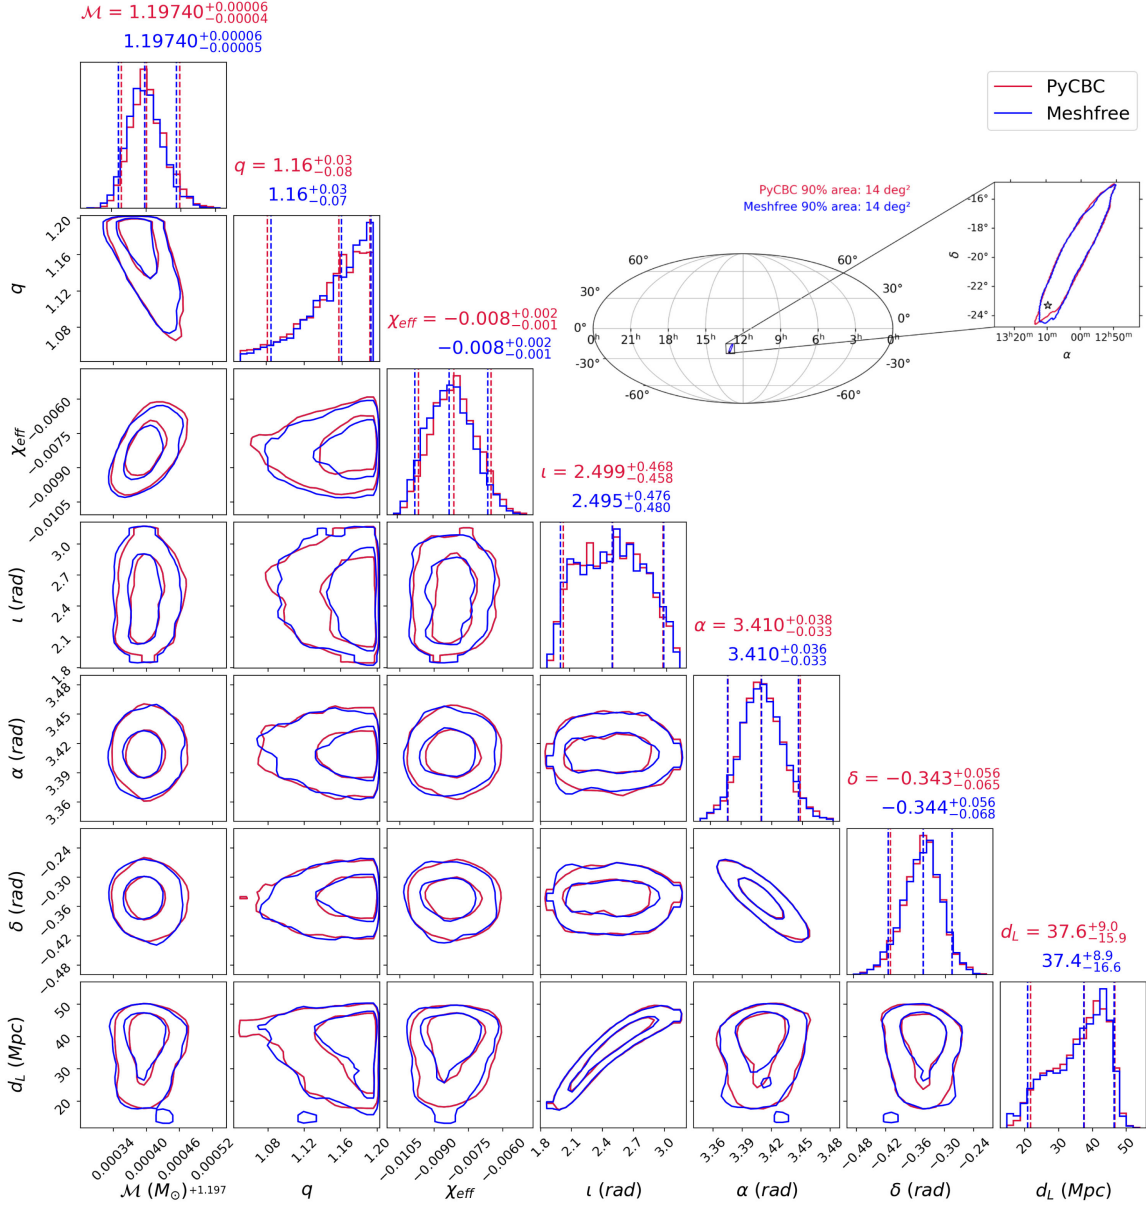| Parameters	Range	Prior distribution
$\mathcal{M}$	$[\mathcal{M}_{\mathrm{cent}} \pm 0.0002]$	$\propto \mathcal{M}$
q	$[q_{\mathrm{cent}} \pm 0.07]$	$\propto [(1+q)/q^3]^{2/5}$
$\chi_{1z,2z}$	$[\chi_{1z\mathrm{cent}} \pm 0.0025]$	Uniform
d_L	$[10, 60]$	Uniform in volume
t_c	$t_{\mathrm{trig}} \pm 0.12$	Uniform
α	$[0, 2\pi]$	Uniform
δ	$\pm \pi/2$	$\sin^{-1}[\mathrm{Uniform}[-1, 1]]$
ι	$[0, \pi]$	Uniform in $\cos \iota$
ψ	$[0, 2\pi]$	Uniform angle

FIG. 3. The corner plot shows the posterior distributions over a ten-dimensional parameter space using both the mesh-free method and the standard PyCBC likelihood. The seismic cutoff frequency is set to 20 Hz, assuming the TaylorF2 waveform model. The estimated distributions from both methods exhibit good agreement, as indicated by the median values in the title of the 1D marginalized posterior plots along the diagonal. The mesh-free likelihood-based PE process was completed in ~3.4 min when performed on 32 cores. A sky map (inset) is also generated for both methods where the star denotes the actual location of the galaxy from which GW170817 is believed to have originated and is in agreement with [74].

The PE run with the mesh-free likelihood (for TaylorF2) completed in ~3.4 min (including the start-up and sampling stages) when run on 32 cores. In a similar analysis with the IMRPhenomD [72] waveform model, we found good agreement between the posterior distributions obtained from the mesh-free and PyCBC likelihood methods, and these results are broadly consistent with the LVK analysis of the same event. In this case, the mesh-free method generated the posterior in ~3.6 min utilizing 32 cores.

As shown in Fig. 3, the posterior distributions estimated using the mesh-free method and standard PyCBC likelihood exhibit good agreement. Additionally, sky maps for GW170817 were generated using the LIGO-SKYMAP [73] utility for two methods, further demonstrating the effectiveness of the mesh-free method in reconstructing source parameters. For the same analysis using 64 cores, the number of likelihood evaluations for mesh-free PE (standard PyCBC PE) is 1396388 (1404588), whereas the number of posterior samples collected during sampling is 3633 (3546).

(a) `TaylorF2` signal model
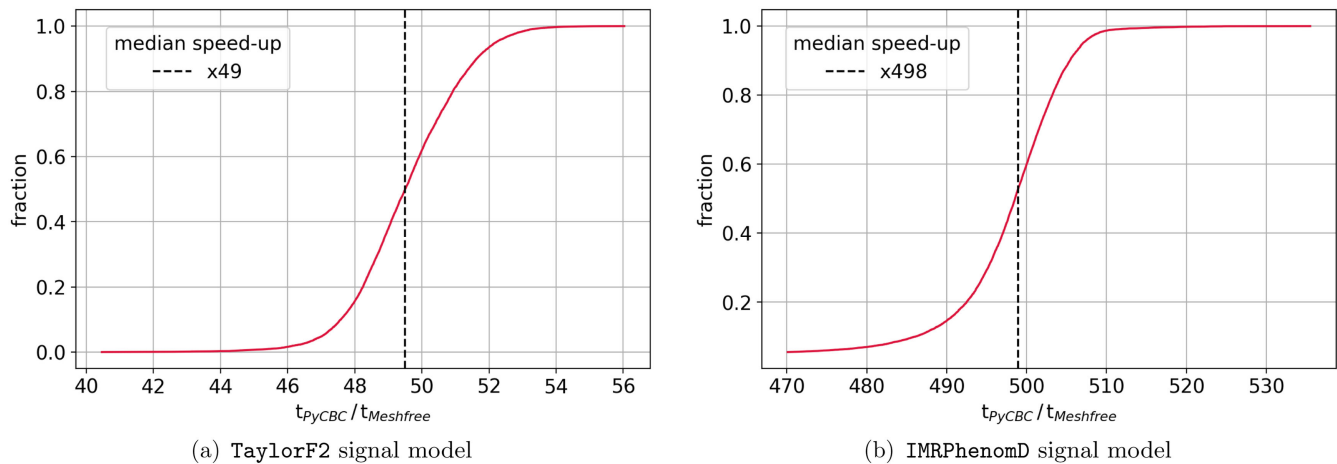

(b) `IMRPhenomD` signal model

FIG. 4.   The plots compare the speed-up of the mesh-free likelihood evaluation with PyCBC likelihood for two waveform models, (a) TaylorF2 and (b) IMRPhenomD. The mesh-free method is around 49 times faster for TaylorF2 and approximately 498 times faster for IMRPhenomD (compared to the standard calculation). The median absolute error is approximately $\mathcal{O}(10^{-2})$ for both cases (see Fig. 9 in Appendix B). In the proposed mesh-free method, one can trade off speed in favor of better accuracy by retaining a larger number of SVD basis vectors [see Eq. (9)] and vice versa. The optimal subjective choice is made from the spectrum of eigenvalues (Fig. 2).

To estimate the speed-up achieved by using the mesh-free likelihood compared to the PyCBC likelihood, we generated 10^4 random query points from the prior distribution specified in Table I. Subsequently, we calculated the log-likelihood at these points using both the mesh-free and PyCBC likelihood methods. The results show a speed-up of ~49 over standard PyCBC likelihood, as shown in Fig. 4. In the case of IMRPhenomD, we observed a relative speed-up of ~498 over standard PyCBC likelihood. This difference in speed-up can be attributed to the difference in time taken to generate the template in the two different waveform models: TaylorF2 waveforms are far less computationally expensive to generate than IMRPhenomD model waveforms. The corresponding error plots are shown in Appendix C.

Furthermore, we investigated the scaling behavior of PE run times (start − up + online) using the mesh-free method with respect to the number of cores employed. For the TaylorF2 analysis of the GW170817 event, we performed the ten-dimensional PE analysis with varying core counts ranging from 8 to 64. Our study indicates that the analysis can be completed in ~2.4 min using 64 cores. Furthermore, we found that the current implementation scales linearly with the number of live points (keeping `dlogz` fixed). It is worth noting that further reductions in estimation times are expected if the number of cores is increased, as the scaling does not saturate at 64 cores. However, as depicted in Fig. 5, the rate at which the estimation times decrease diminishes as the number of cores continues to increase, which may eventually result in run time saturation beyond a certain number of cores.

The total time shown in Fig. 5 includes the time taken by both the start-up and online stages of the mesh-free algorithm, in addition to the time required for the dynamic

nested sampling algorithm to compute the evidence integral up to a desired level of accuracy. To clarify further: (a) The start-up stage of the mesh-free method benefits from the availability of many CPU cores as it can be performed in parallel. (b) Conversely, the online stage is not designed for parallel processing. This is because estimating the likelihood is a relatively fast operation, typically taking $\sim\mathcal{O}(1)$ ms, achieved by evaluating the interpolating functions. This part of the computational cost remains nearly
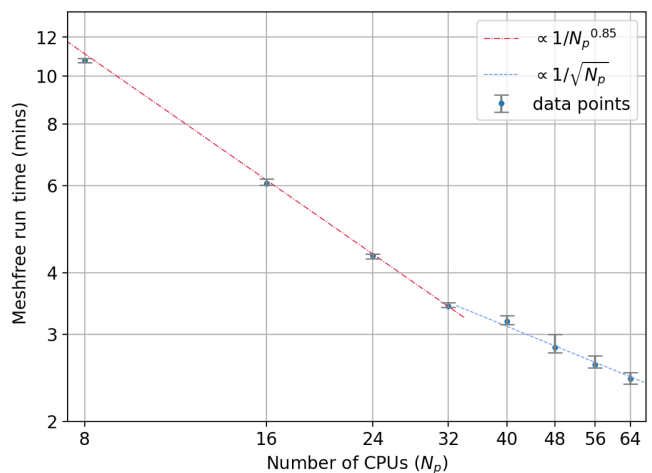

FIG. 5.   The run times of the mesh-free method have been plotted with varying the number of cores. We choose the GW170817 event for this specific example, and the TaylorF2 waveform model is taken for obtaining the posteriors. The plot shows a notable decrease in mesh-free run times as the number of cores increased. We anticipate significant reductions in run times can be achieved with an even higher number of cores, though the gains will eventually saturate as we increase the number of cores.

constant regardless of the specific sample point being processed. As a result, the runtime scaling, as depicted in Fig. 5, effectively combines the time consumed by the DYNESTY sampler and the start-up stage of our method.

B. Simulated data

To assess the performance of our method across a range of SNRs, we generated simulated datasets with different source parameters that mimic BNS systems with network-matched filter SNRs spanning from 10 to 100. In our approach, the number of Gaussian nodes, denoted as N_g, for a given source is intricately tied to its SNR. Specifically, the likelihood profile for events with lower SNRs exhibits a relatively broader spread over the parameter space compared to the more sharply peaked likelihood profiles observed for high SNR events. Consequently, for higher SNR events, we select a larger number of nodes from the Gaussian distribution in contrast to lower SNR events. Thus, we opt for different values of N_g based on the SNR of each source. We will now detail the distribution of source parameters.

The component masses were uniformly drawn from the interval $[1.2, 2.35]M_\odot$, while the magnitude of the dimensionless component spins (aligned with the orbital angular momentum) was uniformly drawn from the interval $[-0.05, 0.05]$. The sources were uniformly distributed in sky location and followed a uniform distribution with d_L ranging from 10 to 200 Mpc. The inclination angle ι of the binary's plane with respect to the line of sight was chosen to be uniformly distributed in $\cos\iota$. The tidal deformability parameter Λ_{tide} of the neutron stars was set to zero for all signals. These signals were simulated using the TaylorF2 waveform approximant, and the noise PSD was chosen from aLIGOZeroDetHighPower [75] for the Livingston and Hanford detectors, and AdvVirgo [76] for the Virgo detector, as implemented in PyCBC.

It is important to recognize that the distribution of injection parameters we have chosen may not hold any astrophysical relevance. These choices have been made solely for the purpose of testing our method across various SNR scenarios, as our node placement algorithm is intricately linked to the SNRs of these events.

At this point, we employ both the mesh-free technique and the relative-binning method (RelBin) for estimating the BNS source parameters. In an ideal scenario, we would select the brute-force likelihood method for comparison, but opting for such an approach might result in extended PE run times. Given that RelBin is a well-established and widely used fast-PE algorithm within the GW community, recognized for its combination of speed and accuracy, we have opted for this alternative over the standard approach.

For all the simulated events, we chose $N = 800$ RBF nodes. However, N_g was empirically chosen depending on
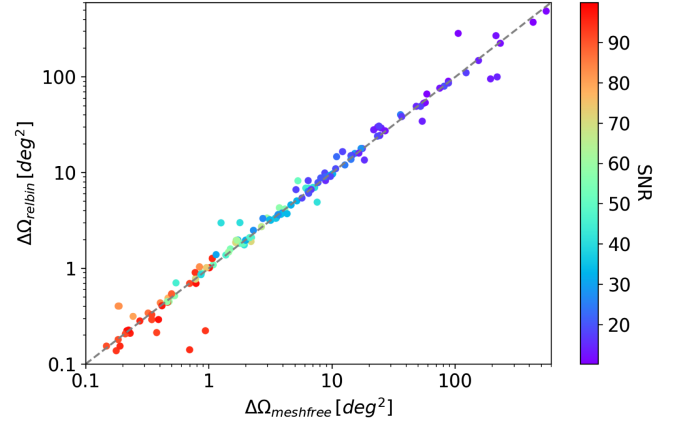

FIG. 6. 90% credible interval (CI) sky area as a function of SNR are shown. We compared the accuracy of the mesh-free method with the RelBin method. The scattered points lie along the positive slope (dashed straight line), implying that the mesh-free method has similar accuracy as RelBin in estimating these parameters, i.e., sky location. It is observed that estimating these parameters with high SNR events is more accurate than with low SNR events. Note that there are some points that are away from the straight line. This is due to the fact that posteriors corresponding to those points are multimodal, giving rise to higher areas for one method in comparison to the other method.

the SNR of the events. Heuristically, we found that, for SNR ≥ 50, $N_g \in [0.4, 0.6]$ was a suitable choice, while for events with SNR ≤ 50, N_g was chosen from $[0, 0.2]$. There were a few exceptions to this prescription as well.

For the RelBin method, we set $\epsilon = 0.000198$, which determines the number of frequency bins in the RelBin method (for more details, refer to [31,32]). This parameter determines the accuracy and run time of the relative-binning method. We choose the reference parameters (or "fiducial point" in the case of RelBin) to be the injection parameters of each simulated event, even though in a realistic scenario, it would have been obtained by an optimization process starting from the best-matched search template.

The prior distributions over the extrinsic parameters $(\alpha, \delta, \iota, \psi, d_L)$ used in the PE are the same as those used for the GW170817 event in the earlier section. As for the intrinsic parameters, we select them according to the following distributions: (i) uniform distribution in component masses (m_1, m_2) such that $\mathcal{M} \in \mathcal{M}_{\text{ref}} \pm 0.0002$ and $q \in q_{\text{ref}} \pm 0.07$,[1] and (ii) uniformly in dimensionless component spins from the interval $\chi_{1z/2z} \in \chi_{1z/2z_{\text{ref}}} \pm 0.0025$, where $\mathcal{M}_{\text{ref}}$, q_{ref}, and $\chi_{1z/2z_{\text{ref}}}$ are chirp mass, mass ratio, and aligned spin components, respectively, corresponding to the reference point found by the optimizing the network SNR starting from the best-matched template. The prior over the time of coalescence t_c is chosen to be uniform

---

[1]In cases where $q_{\text{ref}} - 0.07 < 1$, we sample from $q \in [1, 1.14]$.
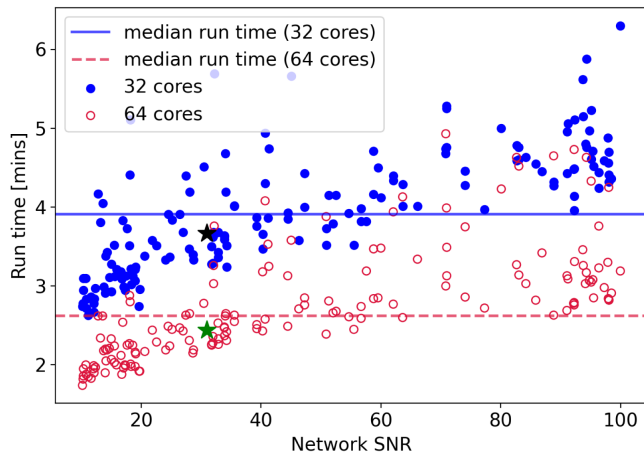
FIG. 7. We present the run times for reconstructing a number of simulated BNS events injected in Gaussian noise (Sec. IV B) using the mesh-free method. Higher SNR injections took longer to reconstruct due to their narrowly peaked likelihood. The analysis utilized 32 (64) cores, represented by blue dots (red circles). The black (green) star shows the GW170817 event's run time on 32 (64) cores for reference. With 64 cores, the median run time decreased to around ∼2.6 min compared to ∼3.9 min with 32 cores. For the RelBin method, using $\epsilon = 0.000198$, the median run time is ∼22 min on 32 cores. The RelBin PE run times crucially depend on the value ϵ, and taking higher values of ϵ can result in lower run times.

from the interval $t_{\mathrm{ref}} \pm 0.12$, where t_{ref} is the geocentric time of the trigger.

As can be seen from the 90% CI sky-area distribution in Fig. 6, the mesh-free method accurately recovers the sky location with a similar level of accuracy as RelBin. The timing results for the simulated signals, shown in Fig. 7, illustrate that the run time increases with the SNR since signals with higher SNRs exhibit a more sharply peaked likelihood function that requires a longer time to accumulate a sufficient number of posterior samples to evaluate the evidence integral up to the desired accuracy.

All the tests were performed on AMD EPYC 7542 CPU@2.90 GHz processors.

V. CONCLUSION AND OUTLOOK

This work extends our previously prescribed mesh-free method [45] to perform parameter estimation on a ten-dimensional parameter space for aligned spin waveform models using data from a multidetector network framework. In the previous version of our method, we demonstrated its effectiveness on a simulated BNS event, as seen in a single detector, where we directly chose the injection parameters as our center for placing uniformly distributed RBF nodes. In this work, we start from the best-matched template parameters and optimize the network SNR to reach a nearby point with a higher network SNR value, which acts as a guide for placing RBF nodes.

This node placement algorithm incorporates a blend of nodes from both multivariate Gaussian and uniform distributions. It enhances the accuracy of likelihood reconstruction (see Appendix A for a study of marginalized posterior profiles on different fractions of random interpolation nodes drawn from a multivariate Gaussian distribution).

We tested our method for several simulated BNS events with varying SNRs and found good agreement between the mesh-free and the relative-binning methods. Further, we have demonstrated that our mesh-free method can rapidly estimate the posteriors of the GW170817 BNS event and efficiently locate the EM counterpart in the sky within ∼2.4 (3.4) min after detection using 64 (32) CPU cores.

These run times quoted above need to be put in the context of other fast-PE algorithms for which we chose the relative-binning method. Note that the time taken by the relative-binning algorithm crucially depends on the choice of the ϵ parameter, which sets the frequency-bin resolution in the heterodyning process. We set $\epsilon = 0.000198$, for which the time taken by the relative-binning method was found to take ∼20 min on 32 cores to analyze the GW170817 event using data from a three-detector (Hanford-Livingston-Virgo) network. This choice of ϵ was taken to match the PE run times claimed by Finstad and Brown [24]. The sampler parameters were kept identical for both methods. Since we do not have the hardware (nor know the parameters) used by Finstad and Brown, this *ad hoc* procedure was used to arrive at this reference time. As such, it would not be prudent to arrive at any definite comparative conclusion based on these timing results. Nevertheless, we will continue to tune the mesh-free method to explore further optimizations for speeding up the PE analysis.

Note that the run time quoted here does not include the time spent in finding the center using optimization, which, when run on a single core, took ∼32 s. However, it can be sped up by running in parallel over multiple CPU cores using multiprocessor versions of the optimization algorithm. For example, see the parallel version of the `scipy.optimize` [77,78] Python routine.

We also used our method for the NSBH event GW200115 [79] (see Fig. 10 in Appendix C) and obtained a reasonably good match of posterior samples between our method and PyCBC standard likelihood.

While our mesh-free framework has the potential to contribute to future LIGO low-latency PE efforts, it is also essential to acknowledge some limitations of the current implementation: First, the RBF mesh-free interpolating functions are seen to be accurate only over a relatively narrow domain in the intrinsic parameter space, which leads to using narrow prior boundaries to carry out the PE analysis. This constraint restricts the flexibility of the mesh-free PE analysis compared to standard

approaches. A potential solution to overcome this issue could be to divide the intrinsic parameter space into smaller "patches" and generate interpolants for each patch independently [80].

In the present implementation, the prior bounds (albeit narrow) were obtained from the best-matched template as explained in Sec. III A 1. Since the PE pipeline is expected to be triggered by an upstream search pipeline, we have used the most significant trigger parameters ($\vec{\lambda}^*$) as the starting point. Using a suitable optimization algorithm, this point ultimately leads us to a nearby reference point ($\vec{\lambda}_{\text{ref}}$), which serves as the center of the hyperrectangle from which the samples are drawn. The dimensions of this hyperrectangle along the different directions are heuristically set to get good accuracy in the likelihood reconstruction (i.e., we do not use the LVK results *a priori* to determine the prior bounds). In the future, we would like to explore an automated way of determining the dimensions of this hyperrectangle from the size of the 90% overlap ellipsoid calculated at $\vec{\lambda}_{\text{ref}}$. Such an idea has been mooted in Pankow *et al.* [81]. However, since the axes of the overlap ellipsoid do not necessarily align with the eigendirections of the data-driven covariance matrix, such an implementation would require careful thought.

Second, it is worth noting that, despite improvements in the placement of nodes compared to our previous work, it still needs investigation for a more generic prescription that can be applied to any CBC system. Incorporating an adaptive strategy [82] to identify significant nodes could enhance interpolation accuracy and strengthen the node placement strategy.

In this analysis, we have focused exclusively on non-precessing signal models (TaylorF2 and IMRPhenomD), with a particular emphasis on the dominant harmonic of the GW signal. These models do not incorporate tidal terms [83,84]. Since the tidal deformability parameters corresponding to the binary components are intrinsic in nature, we need to include them as interpolating parameters in our method, which increases the dimensionality of the intrinsic parameter space to six. Additionally, a greater number of nodes are required to ensure an accurate interpolation of the SVD coefficients and template norm square. We will explore these ideas in future follow-ups of this work.

Although we anticipate that extending our method to include higher-order harmonics should be relatively straightforward, it may result in increased start-up costs. In the case of precessing-spin waveform models [85,86], a mesh-free framework is particularly suitable since the number of RBF nodes does not exponentially increase with the dimensionality of the parameter space. In this regard, the relative-binning algorithm has been recently implemented for gravitational wave parameter estimation with higher-order modes and precession [87]. In the future,

we would like to overcome these limitations by enhancing the node placement algorithm further, expanding the boundaries of the sample space, and incorporating models that account for in plane spins [85].

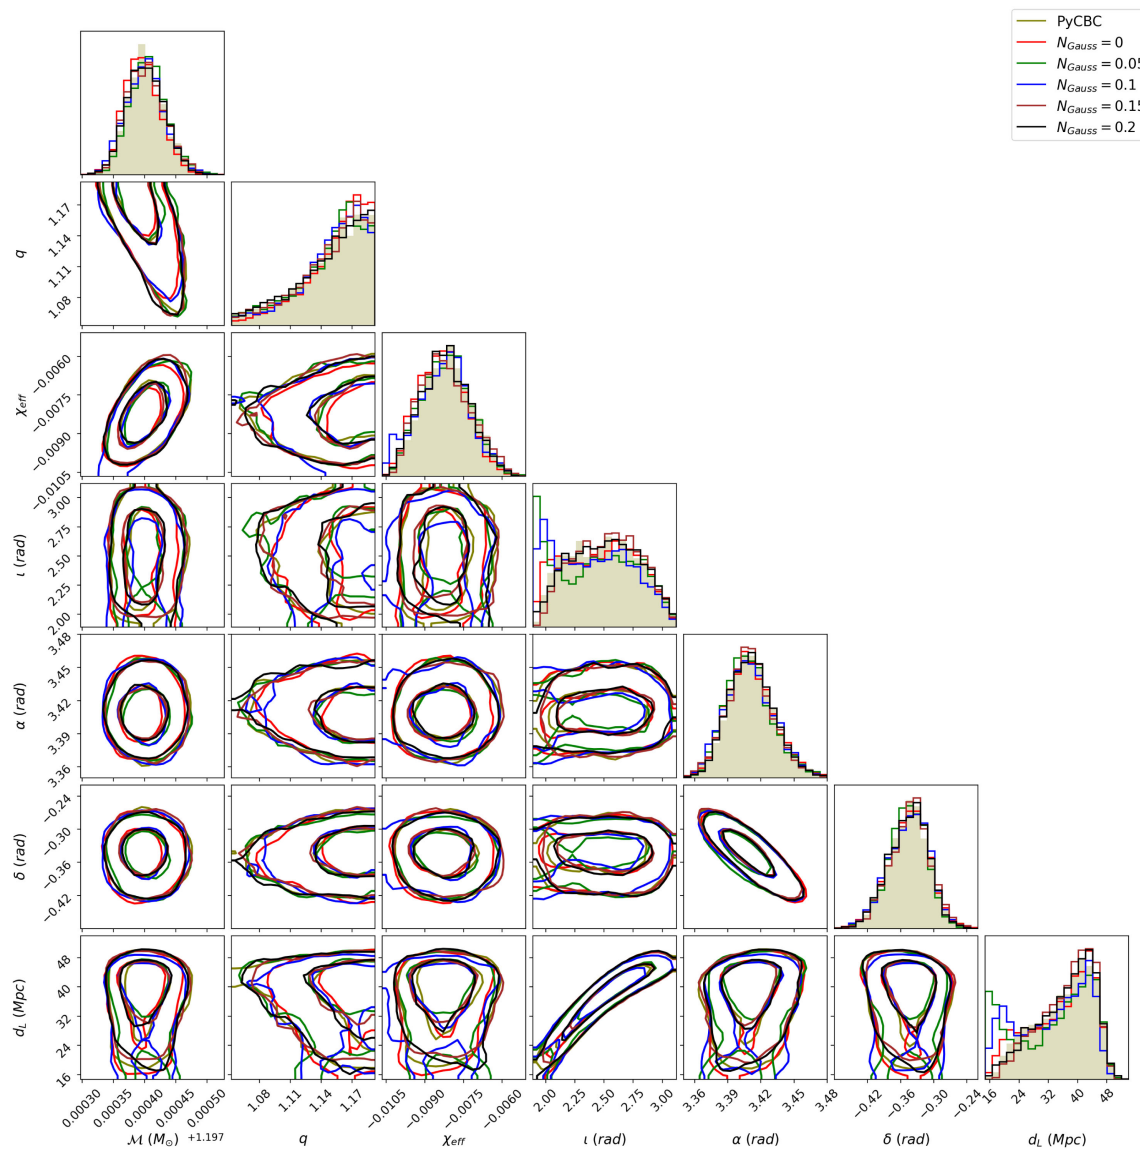Codes used in this analysis are publicly available in the following Github repository [88].

FIG. 8. The effect of choosing different values of N_g (fraction of Gaussian nodes) on the reconstructed parameters.

APPENDIX A: EFFECT OF DIFFERENT CHOICES OF GAUSSIAN DISTRIBUTED NODES ON THE POSTERIORS

The distribution of input nodes plays an important role in the mesh-free PE algorithm as described in Sec. III A 1. One would like to place nodes intelligently so as to provide good coverage near the support of the log-likelihood function, which falls very steeply around its mode. Here, we show the effect of node placement on the posteriors obtained. We choose a certain fraction of RBF nodes (N_g) from a normal (Gaussian) distribution described by the covariance matrix. The remaining nodes are distributed uniformly over the sample space. We find that the overall accuracy increases as we choose a higher fraction of Gaussian distributed nodes (particularly on the ι and d_L

parameters). Although, the estimations of α, δ, and $\mathcal{M}$ are quite robust against different choices of N_g.

APPENDIX B: GW200115: AN NSBH EVENT

We present the results of reconstructing the GW200115 NSBH event using the mesh-free and the standard PyCBC likelihood (for reference). The seismic cutoff frequency is set to 20 Hz, assuming the IMRPhenomD waveform model. We used a 64 s data segment around the GW200115 NSBH event, and PSD was generated using the same data with median-mean estimation. The center of the sample space was chosen as the maximum a posteriori probability values of the PE samples obtained by previous LIGO analysis [90].
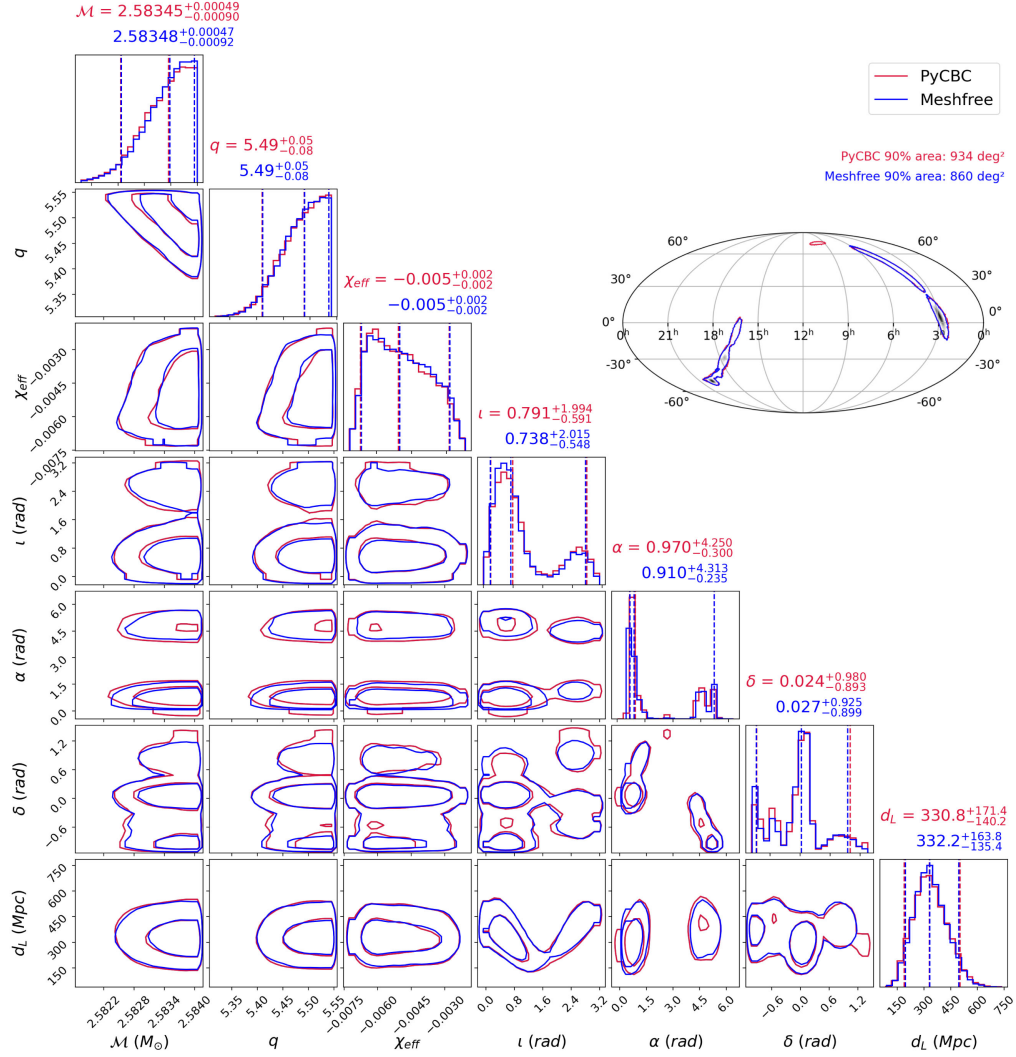
FIG. 9. The corner plot shows the posterior distributions of ten-dimensional parameters using two different methods: the mesh-free method and the standard PyCBC likelihood. The estimated distributions from both methods exhibit good agreement, as indicated by the median values in the marginalized posterior titles. In the inset, the sky maps are also shown, which are in agreement with results published by the LVK Collaboration [91].

The parameters for RBF interpolation are as follows: $N = 1100$, $N_{\text{Gauss}} = 110$, $\nu = 10$, $\phi = \exp(-\epsilon r^2)$, and $\epsilon = 20$. The sampling parameters for DYNESTY nested sampler were the following: (i) nlive = 2500, (ii) walks = 350, (iii) sample = rwalk, and (iv) dlogz = 0.1.

The mesh-free likelihood-based PE process was completed in ∼16 min when performed on 64 cores. The sky maps are broadly in agreement with results published by the LVK Collaboration [91].

APPENDIX C: LIKELIHOOD ERRORS

To assess the accuracy of the mesh-free method to approximate the likelihood, we generated 10^4 random query points from the prior distribution specified in Table I. Then, we calculated the likelihood using both the mesh-free and PyCBC likelihood methods. The median absolute error for both TaylorF2 and IMRPhenomD waveform models was found to be $\mathcal{O}(10^{-2})$.
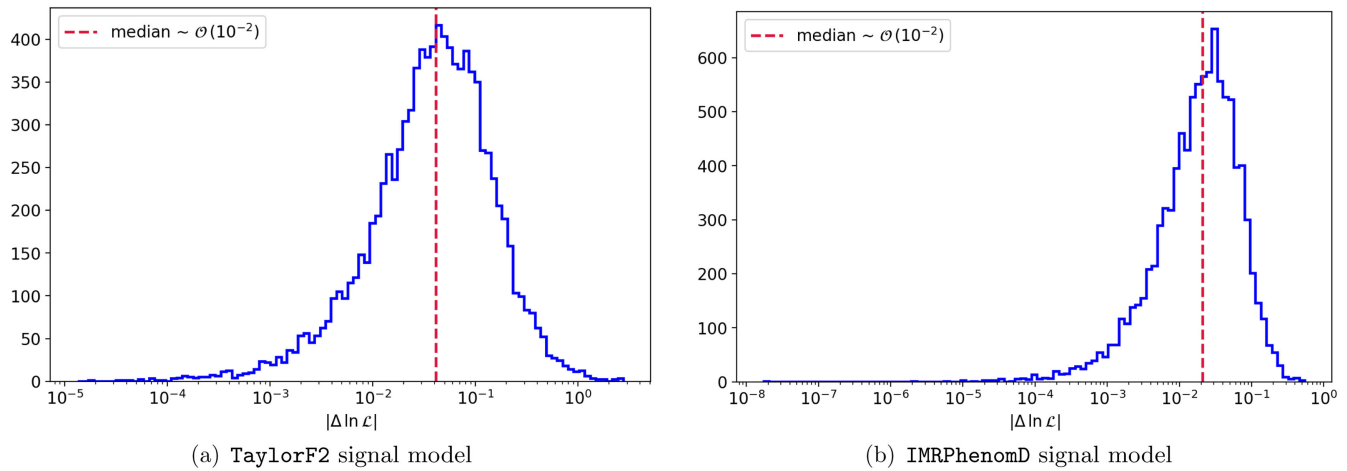
(a) `TaylorF2` signal model

(b) `IMRPhenomD` signal model

FIG. 10.   The likelihood errors ($\Delta \ln \mathcal{L} \equiv \ln \mathcal{L}_{\text{PyCBC}} - \ln \mathcal{L}_{\text{RBF}}$) are shown for both (a) TaylorF2 and (b) IMRPhenomD waveform models (on GW170817 event). The red dashed line represents the median absolute errors ($|\Delta \ln \mathcal{L}|$) that are $\sim \mathcal{O}(10^{-2})$. There are very few points with high errors [$|\cdot| \sim 4$ (2) for TaylorF2 (IMRPhenomD)], which are not shown here.

[1] D. Eichler, M. Livio, T. Piran, and D. N. Schramm, Nucleosynthesis, neutrino bursts and gamma-rays from coalescing neutron stars, Nature (London) **340**, 126 (1989).

[2] J. Goodman, Are gamma-ray bursts optically thick?, Astrophys. J. Lett. **308**, L47 (1986).

[3] B. D. Metzger, Kilonovae, Living Rev. Relativity **1**, 1 (2019).

[4] L.-X. Li and B. Paczyński, Transient events from neutron star mergers, Astrophys. J. **507**, L59 (1998).

[5] I. Gupta, S. Borhanian, A. Dhani, D. Chattopadhyay, R. Kashyap, V. A. Villar, and B. S. Sathyaprakash, Neutron star-black hole mergers in next generation gravitational-wave observatories, Phys. Rev. D **107**, 124007 (2023).

[6] B. P. Abbott, R. Abbott, T. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, V. B. Adya et al., GW170817: Observation of gravitational waves from a binary neutron star inspiral, Phys. Rev. Lett. **119**, 161101 (2017).

[7] G. M. Harry (LIGO Scientific Collaboration), Advanced LIGO: The next generation of gravitational wave detectors, Classical Quantum Gravity **27**, 084006 (2010).

[8] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, V. B. Adya, C. Affeldt, M. Agathos et al., Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA, Living Rev. Relativity **23**, 3 (2020).

[9] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou, A. Allocca, J. Amarni, P. Astone, G. Balestri, G. Ballardin et al., Advanced Virgo: A second-generation interferometric gravitational wave detector, Classical Quantum Gravity **32**, 024001 (2014).

[10] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari et al., Advanced LIGO, Classical Quantum Gravity **32**, 074001 (2015).

[11] D. V. Martynov, E. D. Hall, B. P. Abbott, R. Abbott, T. D. Abbott, C. Adams, R. X. Adhikari, R. A. Anderson, S. B. Anderson, K. Arai et al., Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy, Phys. Rev. D **93**, 112004 (2016).

[12] G. P. Lamb and S. Kobayashi, GRB 170817A as a jet counterpart to gravitational wave trigger GW 170817, Mon. Not. R. Astron. Soc. **478**, 733 (2018).

[13] B. P. Abbott et al. (KAGRA, LIGO Scientific and Virgo Collaborations), Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA, Living Rev. Relativity **23**, 1 (2020).

[14] B. Margalit and B. D. Metzger, The multi-messenger matrix: The future of neutron star merger constraints on the nuclear equation of state, Astrophys. J. Lett. **880**, L15 (2019).

[15] A. K. Lenon, A. H. Nitz, and D. A. Brown, Measuring the eccentricity of GW170817 and GW190425, Mon. Not. R. Astron. Soc. **497**, 1966 (2020).

[16] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, Spin-induced orbital precession and its modulation of the gravitational waveforms from merging binaries, Phys. Rev. D **49**, 6274 (1994).

[17] M. Hannam et al., General-relativistic precession in a black-hole binary, Nature (London) **610**, 652 (2022).

[18] N. E. Wolfe, S. Vitale, and C. Talbot, Too small to fail: Characterizing sub-solar mass black hole mergers with

gravitational waves, J. Cosmol. Astropart. Phys. 11 (2023) 039.

[19] G. Pratten, C. García-Quirós, M. Colleoni, A. Ramos-Buades, H. Estellés, M. Mateu-Lucena, R. Jaume, M. Haney, D. Keitel, J. E. Thompson, and S. Husa, Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes, Phys. Rev. D **103,** 104056 (2021).

[20] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, D. Gerosa, L. C. Stein, L. E. Kidder, and H. P. Pfeiffer, Surrogate models for precessing binary black hole simulations with unequal masses, Phys. Rev. Res. **1,** 033015 (2019).

[21] A. Ramos-Buades, A. Buonanno, H. Estellés, M. Khalil, D. P. Mihaylov, S. Ossokine, L. Pompili, and M. Shiferaw, SEOBNRv5PHM: Next generation of accurate and efficient multipolar precessing-spin effective-one-body waveforms for binary black holes, Phys. Rev. D **108,** 124037 (2023).

[22] T. Islam, J. Roulet, and T. Venumadhav, Factorized parameter estimation for real-time gravitational wave inference, arXiv:2210.16278.

[23] L. P. Singer and L. R. Price, Rapid Bayesian position reconstruction for gravitational-wave transients, Phys. Rev. D **93,** 024013 (2016).

[24] D. Finstad and D. A. Brown, Fast parameter estimation of binary mergers for multimessenger follow-up, Astrophys. J. Lett. **905,** L9 (2020).

[25] C. Chatterjee and L. Wen, Pre-merger sky localization of gravitational waves from binary neutron star mergers using deep learning, Astrophys. J. **959,** 76 (2023).

[26] P. Canizares, S. E. Field, J. R. Gair, and M. Tiglio, Gravitational wave parameter estimation with compressed likelihood evaluations, Phys. Rev. D **87,** 124005 (2013).

[27] P. Canizares, S. E. Field, J. Gair, V. Raymond, R. Smith, and M. Tiglio, Accelerated gravitational wave parameter estimation with reduced order modeling, Phys. Rev. Lett. **114,** 071104 (2015).

[28] R. Smith, S. E. Field, K. Blackburn, C.-J. Haster, M. Pürrer, V. Raymond, and P. Schmidt, Fast and accurate inference on gravitational waves from precessing compact binaries, Phys. Rev. D **94,** 044031 (2016).

[29] S. Morisaki and V. Raymond, Rapid parameter estimation of gravitational waves from binary neutron star coalescence using focused reduced order quadrature, Phys. Rev. D **102,** 104020 (2020).

[30] S. Morisaki, R. Smith, L. Tsukada, S. Sachdev, S. Stevenson, C. Talbot, and A. Zimmerman, Rapid localization and inference on compact binary coalescences with the Advanced LIGO-Virgo-KAGRA gravitational-wave detector network, arXiv:2307.13380.

[31] B. Zackay, L. Dai, and T. Venumadhav, Relative binning and fast likelihood evaluation for gravitational wave parameter estimation, arXiv:1806.08792.

[32] N. J. Cornish, Heterodyned likelihood for rapid gravitational wave parameter inference, Phys. Rev. D **104,** 104054 (2021).

[33] S. Fairhurst, C. Hoy, R. Green, C. Mills, and S. A. Usman, Fast inference of binary merger properties using the information encoded in the gravitational-wave signal, Phys. Rev. D **108,** 082006 (2023).

[34] C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning* (The MIT Press, Cambridge, England, 2006).

[35] J. Lange, R. O'Shaughnessy, and M. Rizzo, Rapid and accurate parameter inference for coalescing, precessing compact binaries, arXiv:1805.10457.

[36] S. Morisaki, Accelerating parameter estimation of gravitational waves from compact binary coalescence using adaptive frequency resolutions, Phys. Rev. D **104,** 044062 (2021).

[37] E. Lee, S. Morisaki, and H. Tagoshi, Mass-spin reparametrization for a rapid parameter estimation of inspiral gravitational-wave signals, Phys. Rev. D **105,** 124057 (2022).

[38] G. Morras, J. F. N. Siles, and J. Garcia-Bellido, Efficient reduced order quadrature construction algorithms for fast gravitational wave inference, Phys. Rev. D **108,** 123025 (2023).

[39] A. J. Chua and M. Vallisneri, Learning Bayesian posteriors with neural networks for gravitational-wave inference, Phys. Rev. Lett. **124,** 041102 (2020).

[40] S. R. Green, C. Simpson, and J. Gair, Gravitational-wave parameter estimation with autoregressive neural network flows, Phys. Rev. D **102,** 104057 (2020).

[41] S. R. Green and J. Gair, Complete parameter inference for GW150914 using deep learning, Mach. Learn. Sci. Tech. **2,** 03LT01 (2021).

[42] H. Gabbard, C. Messenger, I. S. Heng, F. Tonolini, and R. Murray-Smith, Bayesian parameter estimation using conditional variational autoencoders for gravitational-wave astronomy, Nat. Phys. **18,** 112 (2021).

[43] K. W. K. Wong, M. Isi, and T. D. P. Edwards, Fast gravitational wave parameter estimation without compromises, Astrophys. J. **958,** 129 (2023).

[44] R. Legin, M. Isi, K. Wong, A. Adam, L. Perreault-Levasseur, and Y. Hezaveh, Towards unbiased gravitational-wave parameter estimation using score-based likelihood characterization, in *2nd ICML Workshop on Machine Learning for Astrophysics* (ICML, Honolulu, Hawaii, 2023).

[45] L. Pathak, A. Reza, and A. S. Sengupta, Fast likelihood evaluation using meshfree approximations for reconstructing compact binary sources, Phys. Rev. D **108,** 064055 (2023).

[46] A. Nitz, I. Harry, D. Brown, C. M. Biwer, J. Willis, T. D. Canton, C. Capano, T. Dent, L. Pekowsky, A. R. Williamson *et al.*, GWASTRO/PYCBC: v2.0.5 release of PYCBC (2022).

[47] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, EMCEE: The MCMC hammer, Publ. Astron. Soc. Pac. **125,** 306 (2013).

[48] J. Skilling, Nested sampling for general Bayesian computation, Bayesian Anal. **1,** 833 (2006).

[49] J. S. Speagle, DYNESTY: A dynamic nested sampling package for estimating Bayesian posteriors and evidences, Mon. Not. R. Astron. Soc. **493,** 3132 (2020).

[50] S. Koposov, J. Speagle, K. Barbary, G. Ashton, E. Bennett, J. Buchner, C. Scheffler, B. Cook, C. Talbot, J. Guillochon *et al.*, JOSHSPEAGLE/DYNESTY: v2.1.0 (2023).

[51] V. Tiwari, C. Hoy, S. Fairhurst, and D. MacLeod, Fast non-Markovian sampler for estimating gravitational-wave posteriors, Phys. Rev. D **108,** 023001 (2023).

[52] J. Veitch, V. Raymond, B. Farr, W. Farr, P. Graff, S. Vitale, B. Aylott, K. Blackburn, N. Christensen, M. Coughlin *et al.*, Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library, Phys. Rev. D **91**, 042003 (2015).

[53] J. T. Whelan, Notes on the geometry of gravitational wave detection, Technical Report No. LIGO-T1300666-v3, 2013, Notes for lectures at the Banach Centre School of Gravitational Waves, Warsaw, Poland, 2013 July 1–5.

[54] C. Pankow, P. Brady, E. Ochsner, and R. O'Shaughnessy, Novel scheme for rapid parallel parameter estimation of gravitational waves from compact binary coalescences, Phys. Rev. D **92**, 023002 (2015).

[55] E. Thrane and C. Talbot, An introduction to Bayesian inference in gravitational-wave astronomy: Parameter estimation, model selection, and hierarchical models, Publ. Astron. Soc. Aust. **36**, e010 (2019).

[56] K. Cannon, A. Chapman, C. Hanna, D. Keppel, A. C. Searle, and A. J. Weinstein, Singular value decomposition applied to compact binary coalescence gravitational-wave signals, Phys. Rev. D **82**, 044025 (2010).

[57] S. A. Usman, A. H. Nitz, I. W. Harry, C. M. Biwer, D. A. Brown, M. Cabero, C. D. Capano, T. Dal Canton, T. Dent, S. Fairhurst *et al.*, The PyCBC search for gravitational waves from compact binary coalescence, Classical Quantum Gravity **33**, 215004 (2016).

[58] S. Roy, A. S. Sengupta, and P. Ajith, Effectual template banks for upcoming compact binary searches in Advanced-LIGO and Virgo data, Phys. Rev. D **99**, 024048 (2019).

[59] D. Mukherjee, S. Caudill, R. Magee, C. Messick, S. Privitera, S. Sachdev, K. Blackburn, P. Brady, P. Brockill, K. Cannon *et al.*, Template bank for spinning compact binary mergers in the second observation run of Advanced LIGO and the first observation run of Advanced Virgo, Phys. Rev. D **103**, 084047 (2021).

[60] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright *et al.*, SciPy 1.0: Fundamental algorithms for scientific computing in Python, Nat. Methods **17**, 261 (2020).

[61] G. E. Fasshauer, *Meshfree Approximation Methods with Matlab* (World Scientific, Singapore, 2007).

[62] N. Flyer, B. Fornberg, V. Bayona, and G. A. Barnett, On the role of polynomials in RBF-FD approximations: I. Interpolation and accuracy, J. Comput. Phys. **321**, 21 (2016).

[63] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo, SoftwareX **13**, 100658 (2021).

[64] S. V. Dhurandhar and B. S. Sathyaprakash, Choice of filters for the detection of gravitational waves from coalescing binaries. II. Detection in colored noise, Phys. Rev. D **49**, 1707 (1994).

[65] S. Droz, D. J. Knapp, E. Poisson, and B. J. Owen, Gravitational waves from inspiraling compact binaries: Validity of the stationary-phase approximation to the Fourier transform, Phys. Rev. D **59**, 124016 (1999).

[66] G. Faye, S. Marsat, L. Blanchet, and B. R. Iyer, The third and a half-post-Newtonian gravitational wave quadrupole mode for quasi-circular inspiralling compact binaries, Classical Quantum Gravity **29**, 175004 (2012).

[67] L. Blanchet, T. Damour, B. R. Iyer, C. M. Will, and A. G. Wiseman, Gravitational-radiation damping of compact binary systems to second post-Newtonian order, Phys. Rev. Lett. **74**, 3515 (1995).

[68] M. Vallisneri, Use and abuse of the Fisher information matrix in the assessment of gravitational-wave parameter-estimation prospects, Phys. Rev. D **77**, 042001 (2008).

[69] F. Iacovelli, M. Mancarella, S. Foffa, and M. Maggiore, GWFAST: A Fisher information matrix Python code for third-generation gravitational-wave detectors, Astrophys. J. Suppl. Ser. **263**, 2 (2022).

[70] T. Hines, Python package containing the tools necessary for radial basis function (RBF) applications (2015).

[71] H. Rocha, On the selection of the most adequate radial basis function, Appl. Math. Model. **33**, 1573 (2009).

[72] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. J. Forteza, and A. Bohé, Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era, Phys. Rev. D **93**, 044007 (2016).

[73] L. P. Singer and L. R. Price, Rapid Bayesian position reconstruction for gravitational-wave transients, Phys. Rev. D **93**, 024013 (2016).

[74] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya *et al.* (LIGO Scientific and Virgo Collaborations), Properties of the binary neutron star merger GW170817, Phys. Rev. X **9**, 011001 (2019).

[75] L. Barsotti, S. Gras, M. Evans, and P. Fritschel, The updated Advanced LIGO design curve, Technical Report No. LIGO-T1800044-v5, LIGO Scientific Collaboration, 2018.

[76] A. Manzotti and A. Dietz, Prospects for early localization of gravitational-wave signals from compact binary coalescences with advanced detectors, arXiv:1202.4031.

[77] F. Gerber and R. Furrer, optimParallel: An R package providing a parallel version of the L-BFGS-B optimization method, R J. **11**, 352 (2019).

[78] F. Gerber, florafauna/optimParallel-Python: Test zenodo (2020).

[79] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt *et al.*, Observation of gravitational waves from two neutron star–black hole coalescences, Astrophys. J. Lett. **915**, L5 (2021).

[80] S. L. Borne and M. Wende, Domain decomposition methods in scattered data interpolation with conditionally positive definite radial basis functions, Comput. Math. Appl. **77**, 1178 (2019).

[81] C. Pankow, P. Brady, E. Ochsner, and R. O'Shaughnessy, Novel scheme for rapid parallel parameter estimation of gravitational waves from compact binary coalescences, Phys. Rev. D **92**, 023002 (2015).

[82] Q. Zhang, Y. Zhao, and J. Levesley, Adaptive radial basis function interpolation using an error indicator, Numer. Algorithms **76**, 441 (2017).

[83] W. Del Pozzo, T. G. F. Li, M. Agathos, C. Van Den Broeck, and S. Vitale, Demonstrating the feasibility of probing the

neutron star equation of state with second-generation gravitational wave detectors, Phys. Rev. Lett. **111,** 071101 (2013).

[84] T. Dietrich, A. Samajdar, S. Khan, N. K. Johnson-McDaniel, R. Dudi, and W. Tichy, Improving the NRTidal model for binary neutron star systems, Phys. Rev. D **100,** 044003 (2019).

[85] G. Pratten, C. García-Quirós, M. Colleoni, A. Ramos-Buades, H. Estellés, M. Mateu-Lucena, R. Jaume, M. Haney, D. Keitel, J. E. Thompson *et al.*, Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes, Phys. Rev. D **103,** 104056 (2021).

[86] C. García-Quirós, M. Colleoni, S. Husa, H. Estellés, G. Pratten, A. Ramos-Buades, M. Mateu-Lucena, and R. Jaume, Multimode frequency-domain model for the gravitational wave signal from nonprecessing black-hole binaries, Phys. Rev. D **102,** 064002 (2020).

[87] H. Narola, J. Janquart, Q. Meijer, K. Haris, and C. V. D. Broeck, Relative binning for complete gravitational-wave

parameter estimation with higher-order modes and precession, and applications to lensing and third-generation detectors, arXiv:2308.12140.

[88] L. Pathak, github repository of the relevant codes used in this paper (2023), https://github.com/lalit-pathak/prompt_sky_localization_meshfree_pe.

[89] https://gwosc.org/.

[90] R. Abbott, T. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. Adhikari, V. Adya, C. Affeldt, D. Agarwal *et al.*, GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run, Phys. Rev. X **13,** 041039 (2023).

[91] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, D. Agarwal, M. Agathos *et al.*, Observation of gravitational waves from two neutron star–black hole coalescences, Astrophys. J. Lett. **915,** L5 (2021).