

# Fast and faithful interpolation of numerical relativity surrogate waveforms using a meshfree approximation

Lalit Pathak<sup>1,2,\*</sup>, Amit Reza<sup>3,4,5,†</sup> and Anand S. Sengupta<sup>1,‡</sup>

<sup>1</sup>*Indian Institute of Technology Gandhinagar, Palaj Gandhinagar, Gujarat 382055, India*

<sup>2</sup>*Department of Astronomy and Astrophysics, Tata Institute of Fundamental Research, 1, Homi Bhabha Road, Colaba, Mumbai 400005, India*

<sup>3</sup>*Nikhef, Science Park 105, 1098 XG Amsterdam, The Netherlands*

<sup>4</sup>*Institute for Gravitational and Subatomic Physics (GRASP), Utrecht University, Princetonplein 1, 3584 CC Utrecht, The Netherlands*

<sup>5</sup>*Space Research Institute, Austrian Academy of Sciences, Schmiedlstrasse 6, 8042 Graz, Austria*



(Received 28 March 2024; accepted 12 August 2024; published 10 September 2024)

Several theoretical waveform models have been developed over the years to capture the gravitational-wave emission from the dynamical evolution of compact binary systems of neutron stars and black holes. As ground-based detectors improve their sensitivity at low frequencies, the real-time computation of these waveforms can become computationally expensive, exacerbating the steep cost of rapidly reconstructing source parameters using Bayesian methods. This paper describes an efficient numerical algorithm for generating high-fidelity interpolated compact binary waveforms at an arbitrary point in the signal manifold by leveraging computational linear algebra techniques such as singular value decomposition and the meshfree approximation. The results are presented for the time-domain NRHybSur3dq8 inspiral-merger-ringdown waveform model that is fine-tuned to numerical relativity simulations and parametrized by the two component-masses and two aligned spins. For demonstration, we target a specific region of the intrinsic parameter space inspired by the previously inferred parameters of the GW200311\_115853 event, a binary black hole system whose merger was recorded by the network of Advanced LIGO and Virgo detectors during the third observation run. We show that the meshfree interpolated waveforms can be evaluated in  $\sim 2.3$  ms, which is about 38 times faster than its brute-force (frequency-domain tapered) implementation in the PyCBC software package at a median accuracy of  $\sim \mathcal{O}(10^{-5})$ . The algorithm is computationally efficient and scales favorably with an increasing number of dimensions of the parameter space. This technique may find use in rapid parameter estimation and source reconstruction studies.

DOI: [10.1103/PhysRevD.110.064022](https://doi.org/10.1103/PhysRevD.110.064022)

## I. INTRODUCTION

Compact binary systems such as binary black holes (BBHs) or binary neutron stars are one of the most important sources for ground-based gravitational-wave (GW) detectors. The GW signal emitted by such sources can be theoretically modeled in the secular inspiral or postmerger ringdown phase by solving Einstein's field equations. While analytical solutions are not available in the highly nonlinear, so-called “merger” domain of the evolution of these sources, the inspiral and ringdown solutions are often calibrated to a set of numerical relativity (NR) waveforms to construct complete semianalytical waveforms covering the inspiral, merger, and ringdown

(IMR) phases of the evolution of such compact binary sources. Several such waveform models have been developed in both the frequency and time domains, such as the family of frequency-domain IMR waveforms [1–5] and the effective one-body (EOB) family [6,7] of waveforms, which are routinely used for GW data analysis. More recently, NR-based surrogate models have also been developed such as the NRSur7dq2 [8], NRSur7dq4 [9], and NRHybSur3dq8 [10] models, which are among the most accurate waveform models available but are computationally expensive to generate in comparison to their frequency-domain counterparts, such as the IMRPhenomXAS and SEOBNRv4\_ROM models. For reference, NRHybSur3dq8 takes  $\sim 75$  ms to generate a waveform with component masses  $m_{1,2} = (50, 20)M_{\odot}$  and aligned spin parameters  $\chi_{1z,2z} = (0.05, 0.05)$  starting at a seismic cutoff frequency of 15 Hz. For the same set of

\*Contact author: [lalit.pathak@iitgn.ac.in](mailto:lalit.pathak@iitgn.ac.in)

†Contact author: [amit.reza@oeaw.ac.at](mailto:amit.reza@oeaw.ac.at)

‡Contact author: [asengupta@iitgn.ac.in](mailto:asengupta@iitgn.ac.in)

parameters, the time to generate IMRPhenomXAS and SEOBNRv4\_ROM is  $\sim 1.5$  ms and  $\sim 2.4$  ms, respectively.<sup>1</sup>

Our research is inspired by a suggestion by Verma *et al.* [10], who tabled the idea of speeding up the evaluation of the NRHybSur3dq8 time-domain surrogate waveforms by creating a faster frequency-domain variant. In the past, fast frequency-domain surrogates using the technique of model order reduction have been applied to the SEOBNRv4 [6] waveform model, leading to a significant reduction in the computational complexity of evaluating these waveforms. Such an approach could potentially accelerate parameter estimation (PE) for compact binary sources using these accurate waveform models.

There have been several attempts to construct surrogate models of other waveform models. Cannon *et al.* [11] constructed the approximate nonspinning IMR waveforms [12] parametrized by only the two-component masses by first projecting these waveforms over a set of singular-value decomposition (SVD) basis vectors, followed by a grid-based two-dimensional Chebyshev interpolation of the SVD coefficients. Chua *et al.* [13] used artificial neural networks (ANNs) to generate a four-dimensional reduced order model (ROM) of the 2.5 post-Newtonian (PN) frequency-domain TaylorF2 waveform model [14] parametrized by the two component masses and two aligned spins. In their approach, the waveforms are represented as the weighted sums over the reduced basis vectors. The ANNs are then trained to accurately map the GW source parameters into the basis coefficients. Other studies related to building NR surrogate models utilized ANNs [15], Gaussian process regression [16], and deep-learning architectures [17].

In this work, we aim to significantly reduce the computational cost of generating frequency-domain NRHybSur3dq8 waveforms by using a combination of SVD and the meshfree approximation. We focus on illustrating our approach by concentrating on a region of the four-dimensional (two component masses and two aligned spins) parameter space around the inferred source parameters of the GW200311\_115853 [18] event detected during the third observing run of the LVK Collaboration [19]. We construct meshfree interpolants for both the amplitude and phase of these waveforms separately. These interpolants are then used to rapidly evaluate the amplitude and phase of the waveform at any arbitrary query points in the chosen parameter space. These interpolated waveforms can be evaluated in  $\sim 2.3$  ms in comparison to 89 ms ( $\sim 38$  times faster) with the standard (frequency-domain tapered) implementation of time-domain NRHybSur3dq8 in PyCBC [20] with a median error of  $\sim \mathcal{O}(10^{-5})$ . Further, we also perform a Bayesian PE study of a simulated GW200311\_115853-like event coherently injected in Gaussian noise to mimic data from the network

of LIGO and Virgo detectors. The PE runs are carried out using the meshfree interpolated waveforms and take  $\sim 16.4$  minutes to complete on a 64-CPU-core setup and the posterior distributions over the source parameters are found to be broadly consistent with their injected values. In comparison, with a brute-force frequency-domain implementation of tapered time-domain NRHybSur3dq8, the posteriors are generated in  $\sim 5$  h and 50 min (on 64 CPU cores), underscoring the gains achieved with the frequency-domain meshfree waveforms.

The rest of the paper is organized as follows. Section II introduces the NRHybSur3dq8 waveform model and the preprocessing required before we build meshfree interpolants. Section III explains an iterative strategy to find a suitable set of basis vectors using SVD to span the space of amplitude and phase and construct the meshfree interpolants of the resulting SVD coefficients. In Sec. IV we describe various radial basis function (RBF) kernels that can be used to generate meshfree interpolants. Section V demonstrates the results of a PE performed on a simulated BBH event using meshfree interpolants. Finally, we summarize the results in Sec. VI and discuss the limitations of the current implementation, and suggest some ideas for overcoming this limitation in follow-up studies.

## II. NR WAVEFORM MODEL

NRHybSur3dq8 is a time-domain, aligned-spin surrogate model for a “hybridized” nonprecessing NR waveform, which is valid for stellar compact binaries with total masses as low as  $2.5M_{\odot}$ . In this context, the term “hybridized” refers to a combined waveform incorporating both a PN and EOB waveform during the early times, attached smoothly to an NR waveform at late times. The training of this model involves hybridized waveforms derived from 104 NR waveforms, supporting mass ratios  $q \leq 8$  and spin values  $\chi_{1z}; \chi_{2z} \leq 0.8$ . Here,  $\chi_{2z}(\chi_{1z})$  represents the spin of the heavier (lighter) black hole in alignment with the direction of orbital angular momentum. The waveform model also supports spin-weighted spherical harmonic modes with  $l \leq 4$  and the (5, 5) mode but excludes the (4, 0) or (4, 1) modes. In this work, we specifically focus on the dominant, commonly referred to as the “quadrupole” mode (2, 2), present in this waveform model. The GW polarizations, denoted as “plus” ( $h_+$ ) and “cross” ( $h_{\times}$ ), can be concisely represented as a unified complex time series, denoted as  $h = h_+ - jh_{\times}$ . This complex time series can be expressed as a linear combination of spin-weighted harmonic modes  $h_{lm}$  [21,22]. Consequently, GWs propagating along any direction  $(\iota, \phi_0)$  in the binary source’s frame can be expanded in spin-weighted spherical harmonic functions  ${}^sY_{lm}$  with spin weight  $s = -2$  as

$$h(t, \iota, \phi_0) = \sum_{l=2}^{\infty} \sum_{m=-l}^l h_{lm}(t) {}^{-2}Y_{lm}(\iota, \phi_0), \quad (1)$$

<sup>1</sup>The timings quoted here are for the Python implementation in the PyCBC software package.




FIG. 1. NRHybSur3dq8 is generated at a frequency ( $f_{\text{start}} = 10$  Hz) lower than the seismic cutoff frequency ( $f_{\text{low}} = 15$  Hz), which leads to a longer-duration waveform. It is generated for a simulated BBH event whose intrinsic parameters are shown in the title. A time-domain tapering using a half-Kaiser window is performed on the longer waveform to smoothly attenuate the waveform amplitude to zero, which helps to alleviate the Gibbs phenomenon arising due to a jump discontinuity while converting the time-domain waveform to a frequency-domain waveform. The black dashed vertical line represents the epoch at which the amplitudes of the tapered and untapered waveforms become equal (or, equivalently, the value of the half-Kaiser window becomes 1). The sampling frequency is set at 2048 Hz. The waveform amplitude (y axis) is normalized by the maximum value of the strain. The Grey color is for the untapered waveform while red color is for the tapered waveform. The half-Kaiser window is represented by the blue color.

where  $\iota$  represents the inclination angle between the orbital angular momentum of the binary and the line of sight to the detector and  $\phi_0$  corresponds to the initial binary phase. The corresponding frequency-domain waveform for the dominant mode can be written as follows:

$$\tilde{h}_{+/\times} = \mathcal{A}_{+/\times}(f) \exp[j\psi_{+/\times}(f)], \quad (2)$$

where  $\mathcal{A}_{+/\times}(f)$  and  $\psi_{+/\times}(f)$  are the amplitude and phase as functions of frequency, respectively, and  $j$  represents the complex number  $j = \sqrt{-1}$ . Since NRHybSur3dq8 is an aligned-spin waveform model [23–28], the relation  $\tilde{h}_{\times} \propto -j\tilde{h}_{+}$  holds and therefore we only consider  $\tilde{h}_{+}$  and the cross polarization ( $\tilde{h}_{\times}(f)$ ) can be calculated from the above relation. From now on, we drop the subscript “+/\times” from  $\mathcal{A}_{+/\times}$  and  $\psi_{+/\times}$  and simply denote the amplitude and phase by  $\mathcal{A}$  and  $\psi$ , which correspond to plus polarization.

As a preprocessing step, we first calculate the duration of the waveform starting at the seismic cutoff frequency ( $f_{\text{low}}$ ). Subsequently, the starting frequency ( $f_{\text{start}}$ ) of the waveform generation is decreased until we reach twice the duration of the desired waveform. Then, the time-domain NRHybSur3dq8 waveform is generated starting at  $f_{\text{start}}$ , followed by a time-domain tapering to smoothly decrease the amplitude to zero (see Fig. 1). The length of the tapering window is taken as a fraction (0.8 in this case) of the difference between the new and desired duration of the waveform. The longer duration of the waveform makes sure that we do not lose any portion of the waveform in the frequency band to be used in the PE while tapering. We use a Kaiser window [29] (as implemented in PyCBC) for time-domain tapering, which can be expressed as

$$w(n) = I_0 \left( \beta \sqrt{1 - \frac{4n^2}{(M-1)^2}} \right) / I_0(\beta), \quad (3)$$

with  $n \leq \lfloor \frac{M-1}{2} \rfloor$  and where  $I_0$  is the modified zeroth-order Bessel function,  $M$  is the number of points in the output window, and  $\beta$  is the shape parameter, which determines the tradeoff between the main-lobe width and side-lobe level in the Fourier response of the window. As the parameter  $\beta$  becomes large, the main-lobe width increases, while the side-lobe level decreases. For this analysis, we use a half-Kaiser window for tapering (see Fig. 1). The tapering helps to reduce the Gibbs phenomenon while converting the tapered waveform in the frequency domain [see Fig. 2(a)]. Additionally, we also subtract the phase accumulated up to  $f_{\text{low}}$  from the total phase of the frequency-domain waveform (tapered) to make sure the phase at  $f_{\text{low}}$  is zero, which significantly helps in the interpolation of the SVD coefficients corresponding to the phase [see Fig. 2(b)].

In the next section, we lay down an iterative strategy to find suitable basis vectors to span the space of amplitude and phase using SVD. Subsequently, we fit the resulting SVD coefficients for amplitude and phase separately using a linear combination of RBFs and monomials. Since the amplitude and phase are smoothly varying functions of the frequency [see Figs. 2(a) and 2(b)], the corresponding SVD coefficients are expected to exhibit smooth variation over the intrinsic parameter space as well and therefore are suitable for interpolation. Combining the interpolated coefficients at the arbitrary query points within the interpolating region with the corresponding basis vectors gives

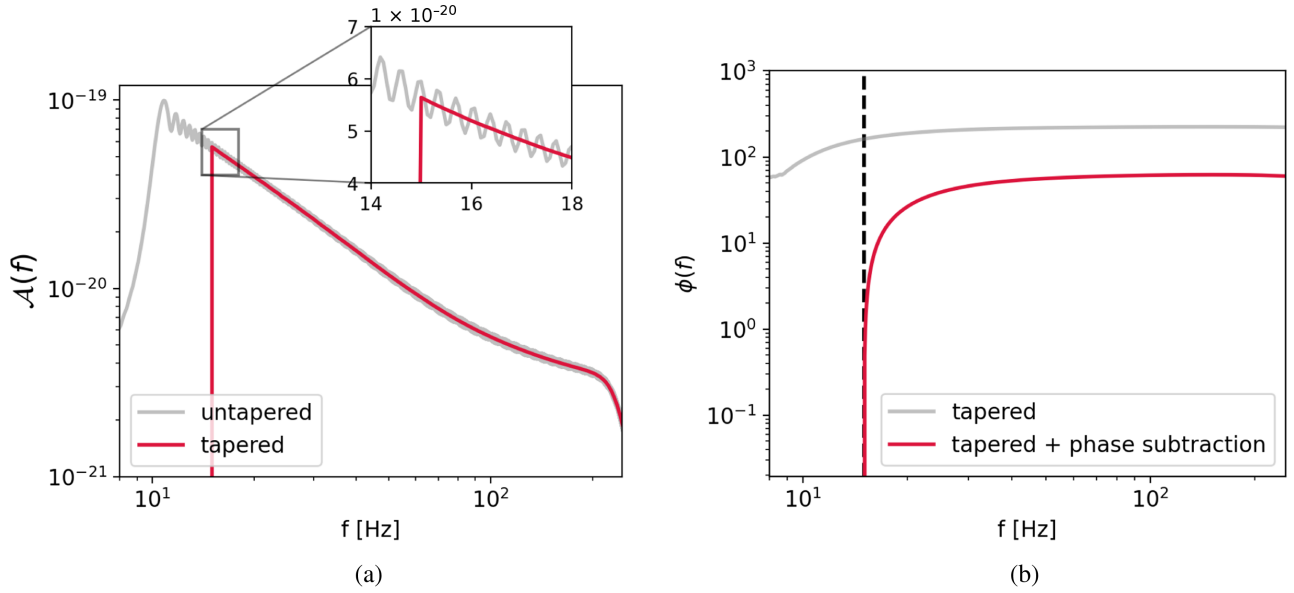


FIG. 2. (a) Amplitude of the untapered waveform (solid grey) is plotted with that of the tapered (solid red) waveform. Note the ringing artifacts appearing due to the Gibbs phenomenon in the untapered waveform. It is almost negligible in the amplitude of the tapered waveform. (b) After converting the tapered waveform into the frequency domain, there is some accumulated phase below 15 Hz, which we subtract from the total phase so that the phase of the frequency-domain waveform is zero at  $f_{\text{low}} = 15$  Hz. The black dashed vertical line represents  $f_{\text{low}} = 15$  Hz.

the interpolated amplitude and phase, and, hence, the interpolated waveform.

### III. GENERATION OF MESHFREE INTERPOLANTS

To construct the interpolants for evaluating the meshfree waveforms, the initial step involves choosing a patch of intrinsic parameter space for which interpolants are to be built. The dimensionality of the patch depends on the dimensionality of the intrinsic parameter space over which interpolation is to be performed. Rather than opting for a single patch with large boundaries, a more effective approach (as suggested in [30]) is to create multiple overlapping patches with smaller boundaries to comprehensively encompass the desired parameter space. For each patch, meshfree interpolants are independently constructed, enhancing the accuracy of individual interpolants. In this analysis, the selection of patches is inspired by a simulated BBH event with GW200311\_115853-like parameters. We use NRHybSur3dq8, an aligned-spin NR surrogate waveform model, which also includes support for subdominant modes. As a proof-of-principle study, this analysis focuses only on the leading-order mode; however, the procedure presented here can be extended to incorporate higher-order modes. The procedure to generate interpolants comprises the following steps.

#### A. Selection of patch(es)

This task involves identifying the patch (or patches) within the parameter space for the construction of waveform

interpolants. A similar approach was taken by Morisaki *et al.* in their work [31] for constructing ROM bases corresponding to different patches of parameter space. While it could be any segment of the pertinent parameter space in general, for this study, we opt for a patch encompassing the injected values of parameters associated with a simulated BBH event, as described earlier. The number of patches should be selected based on the specific research problem, whether it involves fast PE or waveform generation. Our goal is to rapidly obtain posterior distributions, so the number of patches depends on the injection parameters for a simulated signal. For a real GW event, the injection parameters are replaced by the best-matched template. Fewer patches are needed for a PE run because the prior range of parameters is small. The primary criterion is that the patches must encompass the injection parameters, with their width chosen based on expected estimation errors. If the patches do not include the injection parameters, biases in parameter estimation are likely to occur, as in any Bayesian PE analysis. Additionally, the patch width should be selected to generate accurate interpolants, which can be achieved by using multiple patches in the intrinsic parameter space with smaller widths. For waveform generation, a relatively large number of patches is necessary to cover a larger parameter space.

For this work, we define two four-dimensional hyper-rectangular patches with the following ranges:  $\mathcal{M} \in [25, 40]$ ,  $q \in [1, 3]$ ,  $\chi_{1z} \in [-0.1, 0.1]$ , and  $\chi_{2z} \in [-0.05, 0.05]$ . Both patches share common ranges for  $\mathcal{M}$ ,  $q$ , and  $\chi_{2z}$ . However, for the first patch,  $\chi_{1z} \in [-0.1, 0]$ , and for the






FIG. 3. (a) We choose two overlapping patches over  $\chi_{1z}$  to cover the desired range in  $\chi_{1z}$ . A similar procedure can be performed to extend the range of the other parameters ( $\mathcal{M}$ ,  $q$ , etc.). Blue color for the left patch and pink color for the right patch. (b) Final set of nodes after the iterative SVD process terminates for both patches. The nodes found using the “Iterative SVD” procedure tend to cluster near the lower-mass boundaries. Blue points represent the nodes found for patch 1 and orange points represent the nodes found for patch 2.

second patch,  $\chi_{1z} \in [-0.005, 0.1]$ . Note that these two patches overlap in the  $\chi_{1z}$  dimension [see Fig. 3(a)].

### B. Placing grid of RBF nodes using “Iterative SVD” method

This step starts by randomly spraying a minimum (initial) number of nodes across the four-dimensional sample space to start the algorithm and generating frequency-domain waveforms (as specified in Sec. II) at these specified nodes. The minimum number of nodes can be calculated using  $N = \binom{\nu+D}{\nu}$ , where  $\nu$  is the degree of the monomial [see Eq. (5b)] and  $D$  is the dimensionality of the intrinsic parameter space [32,33]. Subsequently,  $\mathcal{A}$  and  $\psi$  are extracted from  $\tilde{h}_+(f)$  for each of the  $N$  waveforms. The obtained amplitudes ( $\mathcal{A}$ ) and phases ( $\psi$ ) are then stacked in a row-wise fashion to construct two matrices,  $\mathbf{A}$  and  $\mathbf{\Psi}$ . These matrices are subjected to SVD, which, in turn, provides the relevant basis for spanning the space defined by these amplitudes and phases,

$$\mathcal{A} = \sum_{\mu=1}^N C_{\mu}^{\mathcal{A}} \vec{u}_{\mu}^{\mathcal{A}} \quad \text{and} \quad \psi = \sum_{\mu=1}^N C_{\mu}^{\psi} \vec{u}_{\mu}^{\psi}. \quad (4)$$

Since by construction the basis vectors ( $\vec{u}_{\mu}^{\mathcal{A}}$ ,  $\vec{u}_{\mu}^{\psi}$ ) are in decreasing order of importance, it turns out that we need to retain only the top few basis vectors (say, up to  $\ell$ ) to accurately reconstruct the amplitude and phase, as illustrated by the singular values spectrum in Fig. 4. Consequently, we limit the construction of interpolants to the SVD coefficients  $C_{\mu}^{\mathcal{A}}$  and  $C_{\mu}^{\psi}$  corresponding to top- $\ell$  basis vectors. Later on, we can choose to retain an even lesser number of basis vectors while reconstructing the amplitude and phase at arbitrary query points. These coefficients, denoted as  $C_{\mu}^{\mathcal{A}}$  and  $C_{\mu}^{\psi}$ , exhibit smooth variations as functions of  $\vec{\lambda}$ , where  $\vec{\lambda} \equiv$

$\{\mathcal{M}, q, \chi_{1z}, \chi_{2z}\}$  represents the interpolating parameter space. These coefficients can be written as a linear combination of RBFs ( $\phi$ ) and monomials ( $p$ ) of specified degree as follows [32]:

$$C_{\mu}^{\mathcal{A}(q)} = \sum_{n=1}^N a_n^{\mathcal{A}} \phi(\|\vec{\lambda}^q - \vec{\lambda}^n\|_2) + \sum_{k=1}^M b_k^{\mathcal{A}} p_k(\vec{\lambda}^q), \quad (5a)$$

$$C_{\mu}^{\psi(q)} = \sum_{n=1}^N a_n^{\psi} \phi(\|\vec{\lambda}^q - \vec{\lambda}^n\|_2) + \sum_{k=1}^M b_k^{\psi} p_k(\vec{\lambda}^q). \quad (5b)$$




FIG. 4. Spectrum of singular values corresponding to amplitude  $\mathcal{A}$  and phase  $\varphi$  [see Eq. (4)], normalized to the maximum value. Note that singular values (relative to the maximum) drop to  $\sim \mathcal{O}(10^{-6})$  within the top 70 basis vectors. Blue and Olive colors represent the spectrum of amplitude and phase respectively while the gray dashed line represents the top-70 basis vectors.

The addition of these monomial terms enhances the accuracy of RBF interpolation, as shown in [34]. From the above equations, it is clear that we need  $(N + M)$  equations to uniquely solve for the  $(N + M)$  RBF and monomial coefficients ( $a^{A/\psi}$  and  $b^{A/\psi}$ ). The SVD coefficients ( $C_\mu^{A(q)}$  and  $C_\mu^{\psi(q)}$ ) are known at the  $N$  RBF nodes.  $M$  additional conditions are imposed by demanding  $\sum_{k=1}^M a_k^{A/\psi} p_k(\lambda^q) = 0$ , which leads to a system of  $(N + M)$  linear equations that can be used to uniquely determine the RBF and monomial coefficients. Using these RBF coefficients, we can quickly evaluate the SVD coefficients  $C_\mu^{A(q)}$  and  $C_\mu^{\psi(q)}$  (corresponding to the  $\mu$ th SVD basis vector) at an arbitrary query point within the interpolating parameter space, and then combine it with corresponding basis vectors ( $\vec{u}^A$  and  $\vec{u}^\psi$ ) to get back the interpolated amplitude  $\mathcal{A}$  and phase  $\psi$ . Then, we combine both the amplitude and phase in accordance with Eq. (2) to get the interpolated frequency-domain waveform. Once we have the interpolants ready, a set of  $10^3$  query points is generated within the interpolating patch, and corresponding true frequency-domain waveforms and meshfree interpolated waveforms are generated at these query points. Subsequently, the mismatches between the true and interpolated waveforms are calculated. A specified number of query points (ten for this work) exhibiting the worst mismatch are added back into the original set of initial nodes. The process then reverts to Sec. III B and continues until either a limit of maximum nodes is reached or the maximum mismatch at the current iteration falls below a predefined threshold (e.g.,  $\leq 10^{-4}$ ) on maximum mismatch. Note that the mismatches are calculated between meshfree waveforms and waveforms first generated in the time domain at  $f_{\text{start}} = 10$  Hz, and then tapered and converted into the frequency domain followed by phase subtraction, as mentioned in Sec. II. We define the match between the two waveforms  $h_1$  and  $h_2$  as

$$\mathcal{O}(h_1, h_2) = \frac{\langle h_1 | h_2 \rangle}{\sqrt{\langle h_1 | h_1 \rangle \langle h_2 | h_2 \rangle}}, \quad (6)$$

where

$$\langle h_1 | h_2 \rangle = \int_{f_{\text{low}}}^{f_{\text{high}}} \frac{\tilde{h}_1(f) \tilde{h}_2^*(f)}{S_n(f)} df \quad (7)$$

which defines a noise-weighted inner product, where  $S_n$  is the one-sided noise power spectral density and the asterisk (\*) denotes complex conjugation. Finally, the mismatch is defined as

$$\text{Mismatch} = 1 - \mathcal{O}(h_1, h_2), \quad (8)$$

where the match  $\mathcal{O}(h_1, h_2)$  is maximized over time and phase shift between the two waveforms.

We call this process ‘‘Iterative SVD’’ because it progressively refines the meshfree waveform approximation by reinserting the query points with the worst mismatch back into the initial set of nodes, followed by evaluation of the true frequency-domain waveforms at the appended set of nodes.

Currently, with each iteration of the ‘‘iterative SVD’’ node placement algorithm, as additional RBF nodes are introduced, the data matrices  $\mathbf{A}$  and  $\mathbf{\Psi}$  are expanded by incorporating more rows (one for each newly added node) followed by an SVD conducted on these expanded data matrices. This approach is practical for handling the decomposition of the small-sized data matrices. However, one can update the basis vectors and singular values at each iteration with more sophisticated algorithms [35,36] that update the left and right subspaces and singular values from the existing SVD decomposition of a dense matrix as new rows are added to it. For updating the subspaces and singular values, Brand *et al.* [35] developed an algorithm for computing a thin SVD of a data matrix of size  $(m \times n)$  in a single pass with linear time complexity  $\mathcal{O}(mnr)$ , where  $r \leq \sqrt{\min(m, n)}$ . The authors proposed fast and memory-efficient sequential algorithms for tracking the singular values and subspaces, initialized with a general identity included in the existing decomposing of the data matrix. Adding a new column would modify the identity representation and provide rank-one updates via the modified Gram-Schmidt algorithm.

Note that if the required number of RBF nodes is small, changing the strategy for computing the subspace and coefficients using the ‘‘updated SVD’’ approach would not significantly reduce the overall computational cost. On the other hand, for a large number of nodes, the incorporation of such sophisticated linear algebra algorithms in our proposed framework would be beneficial. We will explore this possibility in detail in future work.

Next, the SVD is performed for the new amplitude and phase matrices. This is a similar approach to that used in previous studies for building ROMs of a number of waveform approximants [31,37–40]. In each iteration, SVD is employed to identify improved basis vectors for  $\mathcal{A}$  and  $\psi$ . As shown in Fig. 5, the maximum mismatch calculated between the worst meshfree waveform and the corresponding true waveform decreases as we add the query points corresponding to the worst meshfree waveforms at each iteration into the set of nodes used to generate the interpolants. The calculation of mismatches assumes a aLIGOZeroDetHighPower [41] power spectral density (PSD) as implemented in PyCBC. We repeat the same ‘‘Iterative SVD’’ procedure for the other interpolating patches (if needed). For the analysis in this paper, we choose two patches (see Sec. V). As is evident from Fig. 3(b), RBF nodes (for both patches) found by the ‘‘Iterative SVD’’ tend to cluster more toward the low-mass boundaries, implying that the points near the lower-mass boundaries contribute the




FIG. 5. Mismatch between the worst approximated meshfree waveform and true waveform at each iteration. As is evident from the figure, with each iteration, the meshfree approximation becomes better. The mismatch between the worst approximated meshfree waveform and the true waveform at each iteration is represented by red color.

most in terms of the suitable basis vectors for each patch. Similar patterns are also seen in the work focusing on building a surrogate of PN waveforms [10] to find the desired training data set of the parameters and building ROM using ANNs [42], using a greedy strategy.

Note that, initially, the maximum mismatch at each iteration is decreasing rapidly. After  $\sim 25$  iterations, the rate of decrease in the mismatch reduces despite adding the ten worst points at each iteration. This implies that it is close to having the minimum number of nodes required to satisfy the maximum mismatch threshold criteria mentioned earlier. After  $\sim 48$  iterations, the maximum mismatch falls below  $10^{-4}$ , and the algorithm stops. However, to confirm whether it has indeed crossed this criterion (since there are fluctuations), we can test the accuracy of interpolants again by spraying a different set of random query points within the interpolating parameter space and confirming that the maximum mismatch is still less than  $10^{-4}$ . Otherwise, we can continue the algorithm until this threshold criterion is met by a few sets of random query points.

#### IV. CHOICE OF RBF KERNEL

There are different RBF kernels that can be used as a basis [see Eq. (5b)] for the meshfree approximation. In this work, we try the following three different RBF kernels (see Table I): (i) Inverse multiquadric (“imq”), (ii) Inverse Quadric (“iq”), and (iii) Gaussian (“ga”). As mentioned in Sec. III, to terminate the iterative SVD, we set a threshold of either a maximum number of nodes (3000 here) or a maximum mismatch of  $10^{-4}$  or less at a given iteration. The interpolant generation is stopped whenever any of the above-stated conditions is satisfied. In the case of

TABLE I. Different RBF kernels.

RBF kernel type	$\phi(r)$	$\epsilon$	
		Patch 1	Patch 2
inverse multiquadric (“imq”)	$1/\sqrt{1+(er)^2}$	0.4016	0.4010
Gaussian (“ga”)	$\exp[-(er)^2]$	0.3172	0.2909
inverse quadric (“iq”)	$1/(1+(er)^2)$	0.4464	0.4378

“ga” (“iq”) kernels, the final number of nodes for patch 1 and patch 2 is 1530(900) and 1620(590), respectively, when the maximum mismatch at an iteration falls below the threshold of  $10^{-4}$ . However, in the case of the “imq” RBF kernel, the final number of nodes for patch 1 was 680, while for patch 2 it was 770 for the same threshold. A natural question arises: what is so special about the “imq”? A possible explanation is the spread of these RBF kernels. As shown in Fig. 6, the “ga,” “imq,” and “iq” kernels have similar profiles near the center ( $r = 0$ ). However, as we go away from the center, the “ga” and “iq” kernels fall very rapidly in comparison to the “imq” kernel. This implies that the “imq” kernel has a significant overlap with the “imq” kernel centered at the other nearby nodes (centers) in comparison with the other two RBF kernels, and therefore it explains the lesser number of nodes required to fall below the maximum mismatch threshold at a given iteration for “imq” in comparison with a relatively higher number of nodes for the other RBF kernels. For the next section, we choose “imq” as the RBF kernel for generating meshfree interpolants.




FIG. 6. Profiles of different RBF kernels. In the legend, numbers parentheses are the final number of nodes required (for patch 1) to fall below the maximum mismatch threshold of  $10^{-4}$  in the iterative SVD process. For the “imq” kernel, we require the least number of nodes to achieve the threshold. The profiles of imq, iq, and ga are represented by blue, olive, and red colors respectively.

TABLE II. Injected parameters of the simulated BBH event. We take  $q = m_1/m_2$  where  $m_1 \geq m_2$ . These injection parameters mimic the inferred source parameters of the GW200311\_115853 event detected in the O3 science run of the Advanced LIGO and Virgo detectors.

Parameters	Injected values
$\mathcal{M}$	$32.72M_\odot$
$q$	1.22
$\chi_{\text{eff}}$	-0.0091
$d_L$	2155 Mpc
$t_c$	1267963151.3
$\alpha$	0.036 rad
$\delta$	-0.134 rad
$\iota$	0.518 rad
$\psi$	0 rad

## V. RESULTS

As previously mentioned, there are two overlapping patches in  $\chi_{1z}$ , and we independently construct interpolants for each of the patches. In the first patch where  $\chi_{1z} \in [-0.1, 0]$ , we initiate with  $N = (\nu^+ d)$  nodes, with  $\nu$  representing the degree of the monomial terms in Eq. (5b) and  $d$  denoting the dimensionality of the parameter space (here,  $\nu = 6$  and  $d = 4$  for the four-dimensional parameter space interpolation). The “imq” RBF kernel is selected, where  $\epsilon$  is the shape factor influencing the RBFs’ spread at the nodes, determined through the leave-one-out cross validation procedure [43] (see Table I for values of  $\epsilon$ ). In the second patch covering  $\chi_{1z} \in [-0.005, 0.1]$ , the RBF parameters chosen for the first patch are retained, with the exception of a different  $\epsilon$  value for this patch. The seismic cutoff frequency ( $f_{\text{low}}$ ) is set at 15 Hz, while  $f_{\text{start}}$  is equal to 10 Hz.

Once the interpolants for both patches are ready (see Sec. III), we can use them to generate the interpolated waveforms at the arbitrary query points within the interpolating sample space. Note that the generation of interpolants is a one-time offline process and it can be

TABLE III. Prior parameter space over the ten-dimensional parameter space  $\tilde{\Lambda}$ .

Parameters	Range	Prior distribution
$\mathcal{M}$	[25, 40]	$\propto \mathcal{M}$
$q$	[1, 3]	$\propto [(1+q)/q^3]^{2/5}$
$\chi_{1z}$	[-0.1, 0.1]	Uniform
$\chi_{2z}$	[-0.05, 0.05]	Uniform
$V_{\text{com}}$	$[5 \times 10^3, 1 \times 10^{11}]$	Uniform
$t_c$	$t_{\text{trig}} \pm 0.12$	Uniform
$\alpha$	$[0, 2\pi]$	Uniform
$\delta$	$\pm\pi/2$	$\sin^{-1}[\text{Uniform}[-1, 1]]$
$\iota$	$[0, \pi]$	Uniform in $\cos \iota$
$\psi$	$[0, 2\pi]$	Uniform angle




FIG. 7. Mismatches calculated between interpolated and true waveforms at the posterior samples are shown. For comparison, we also show the mismatches for meshfree waveforms using “ga” and “iq” RBF kernels. The median mismatch (for all kernels) is  $\mathcal{O}(10^{-5})$ . Note that the accuracy achieved with “ga” is a little better than with the other kernels due to the high number of nodes that were used to construct “ga” kernel-based meshfree interpolants. However, we do not find any significant difference in the corresponding posterior distributions. The mismatches between the interpolated and true waveforms for imq, iq, and ga are represented by blue, olive, and red colors respectively.




FIG. 8. Absolute relative errors in the likelihood as a function of true  $\log \mathcal{L}$ . The absolute relative error is defined as  $|(\log \mathcal{L}_{\text{mf}} - \log \mathcal{L}_{\text{brute}}) / \log \mathcal{L}_{\text{brute}}|$ .  $\log \mathcal{L}_{\text{mf}}$  is calculated using the interpolated waveforms, while  $\log \mathcal{L}_{\text{brute}}$  is calculated using the true waveforms. These likelihoods are calculated at the posterior samples obtained from the PE of the simulated BBH event. The red star represents the maximum likelihood posterior sample. Note that the majority of relative errors (especially near the peak) are well within 1%, demonstrating the accuracy of the interpolants. The blue points represent the relative error in the likelihood while the red star represents the likelihood corresponding to the maximum likelihood posterior sample.



completed well in advance before the parameter estimation of the GW event is initiated. While multiple patches are necessary to adequately cover the desired parameter space ranges, the generation of interpolants is highly parallelizable and can benefit from a multicore setup to expedite this stage. In the specific case of this analysis, the generation of interpolants for each patch is accomplished in  $\sim 20$  minutes using a CPU setup with 32 cores each. We use the publically available RBF Python package [44] to generate RBF interpolants.

To assess the accuracy of the interpolated waveforms, we generate 20 sets of  $10^3$  query points each, randomly distributed in the intrinsic parameter space within the interpolating region (see Table III). We then generate both the true and interpolated waveforms at these query points and compute the mismatches between them, assuming a flat PSD. The median mismatch is found to be  $\mathcal{O}(10^{-5})$ , which shows the high level of accuracy of the meshfree waveforms in approximating the true waveforms across the interpolating sample space.

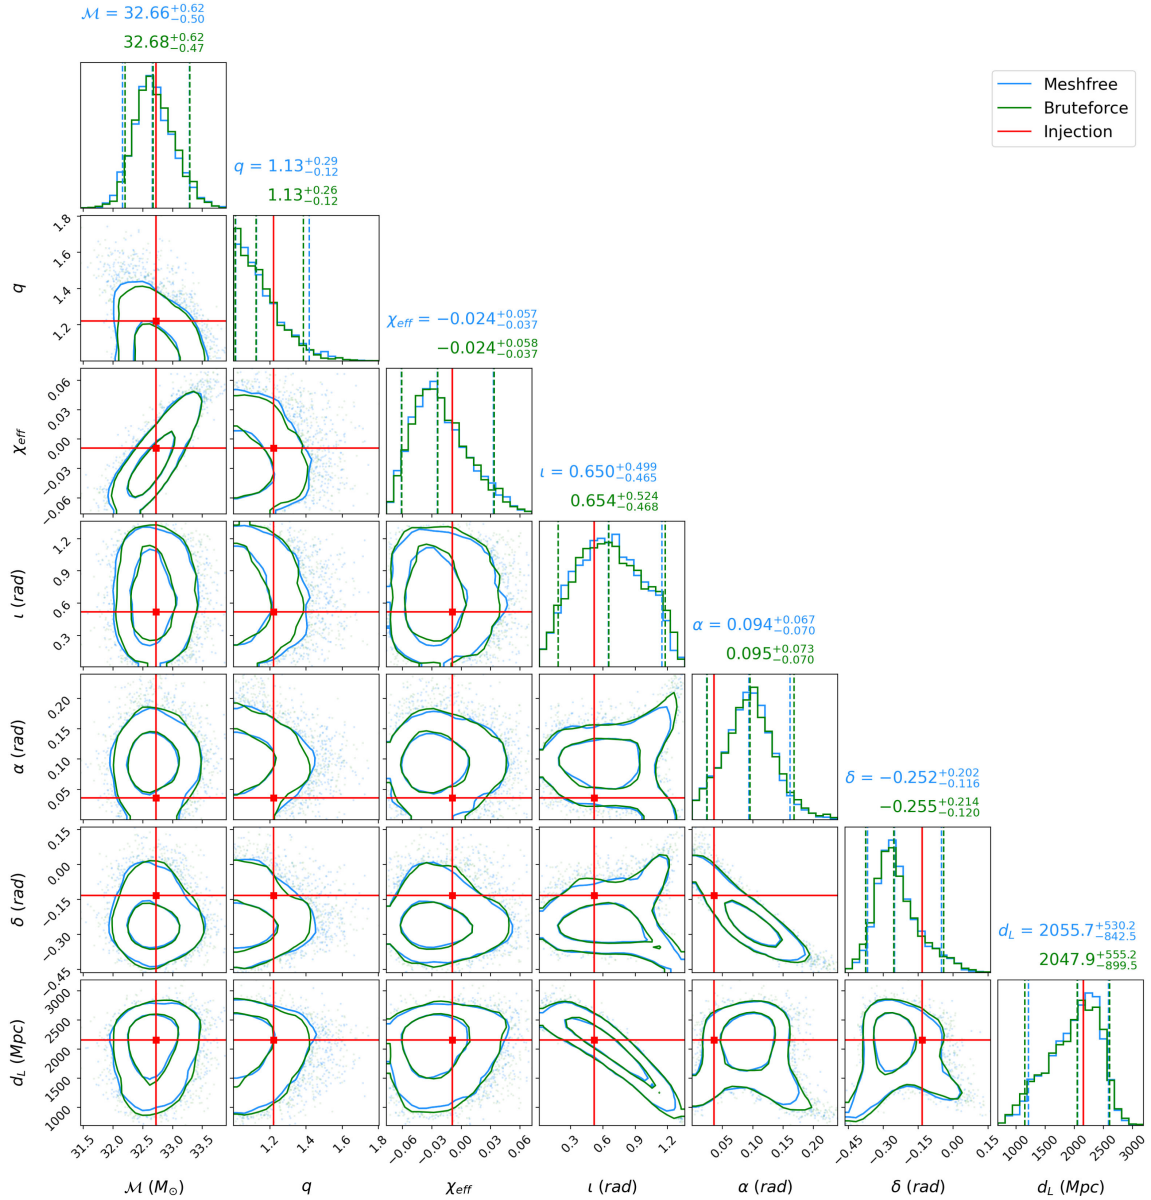


FIG. 9. Corner plot of the posterior distributions (with 50% and 90% CIs and contours) obtained by the meshfree interpolated waveforms for a simulated BBH event. The vertical red line represents the injected values. It takes  $\sim 16.4$  min to complete on 64 CPU cores. In comparison, the PE run with the direct evaluation of the frequency-domain version of the tapered time-domain NRHybSur3dq8 waveform model takes  $\sim 5$  h and 50 min on the same number of CPU cores and the resulting posterior distributions are broadly consistent with those obtained using meshfree interpolated waveforms. The posteriors generated using the Bruteforce waveform are represented by olive color while those generated by the meshfree waveforms are in blue color. The red vertical and horizontal lines represent the injected value (injection).

In terms of the computational speedups, the meshfree waveforms can be evaluated in  $\sim 2.3$  ms in comparison with  $\sim 89$  ms (using the frequency-domain tapered version of the NRHybSur3dq8 time-domain model), implying a speedup factor of  $\sim 38$ . Note that these timings are for the Python implementation of these waveforms. A “C implementation” of the true waveforms can be evaluated in  $\sim 10$  ms, as shown in [10] (about 4.3 times slower than the meshfree waveform’s Python implementation). Since these numbers are also dependent on the hardware configurations of the machines on which these speed tests are performed, a fair comparison of meshfree speedups with their C implementation is not possible until we make a C version of meshfree waveforms. Finally, to test the accuracy of the meshfree model in the context of parameter estimation, we use the SEOBNRv4\_ROM [6] waveform model to inject a 16-second long simulated BBH event (see Table II) into the Gaussian noise with PSDs taken from aLIGOZeroDetHighPower [41] for both the LIGO-Livingston and LIGO-Hanford detectors, and AdvVirgo [45] for the Virgo detector with a network matched-filter signal-to-noise ratio  $\sim 18$ . The seismic cutoff frequency ( $f_{\text{low}}$ ) is set at 15 Hz and the high cutoff frequency is equal to the ringdown frequency of the lowest component masses within the chosen range of chirp masses and mass ratios. A sampling frequency of 2048 Hz is considered. We employ the Dynesty sampler [46,47], a Python implementation of the nested sampling algorithm, to sample the posterior distribution. The prior distributions and the boundaries of the parameters (to be estimated) are shown in Table III. We choose the following sampler settings for DYNESTY: `nLive = 500`, `nWalks = 500`, `dlogz = 0.1`, `sample = “rwalk.”` Here the parameter `nLive` represents the number of live points, which determines the resolution of the sampled posterior distribution. Smaller values of `nLive` might give rise to a poorly sampled distribution (and hence evidence) with much faster convergence. Instead, taking a larger value can give us a finely sampled distribution at the expense of lower convergence speeds. `nWalks` specifies the minimum number of steps required before a new live point is proposed, which replaces the live point with the lowest likelihood in the nested sampling. `sample` indicates the chosen approach for generating samples, and `dlogz` characterizes the remaining prior volume’s contribution to the total evidence. In this PE study, the sampling terminates once `dlogz` reaches a threshold of 0.1. For a more comprehensive understanding of Dynesty’s nested sampling algorithm and its practical implementation, one can refer to [46,48].

As is evident from Fig. 9, posterior distributions of various binary parameters of the simulated BBH event contain the injected values well within the 90% CI. This PE analysis takes  $\sim 16.4$  minutes on a 64-CPU-core setup. Figure 7 also shows the probability distribution function of

the mismatches of the interpolated waveforms with the true waveforms generated at the posterior samples obtained from the PE and the median mismatch is  $\sim \mathcal{O}(10^{-5})$ , demonstrating the good accuracy of the meshfree waveforms. To quantify the effect of these approximate waveforms on the likelihood calculation, we also evaluate the relative errors in the likelihood, as shown in Fig. 8, and find a median absolute relative error of  $\sim \mathcal{O}(10^{-3})$ , further showing the effectiveness of the meshfree waveforms.

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

## VI. CONCLUSION AND FUTURE OUTLOOK

In this work, we constructed a fast, high-fidelity frequency-domain model for the time-domain NRHybSur3dq8 NR waveform model. We showed that such interpolated waveforms can be evaluated in  $\sim 2.3$  ms,  $\sim 38$  times faster than the frequency-domain tapered version of the time-domain implementation of this waveform model in PyCBC. We constructed the waveform interpolants by first selecting a patch of the parameter space. Then, we sprayed randomly scattered nodes in the sample space and evaluated  $\mathcal{A}$  and  $\varphi$  from the frequency-domain waveform corresponding to NRHybSur3dq8 at those nodes. Subsequently, we found a suitable basis spanning the space of both amplitude and phase by performing the SVD of the amplitude and phase matrices resulting from stacking the amplitudes and phases at the nodes. The resulting SVD coefficients corresponding to the amplitude and phase were fit with a linear combination of RBFs and monomials (“interpolants”), followed by the mismatched calculation of the interpolated waveforms with the true waveforms at arbitrary query points in the selected patch of parameter space. The ten worst query points in terms of mismatch were added back into the set of input nodes, and the whole process (called “Iterative SVD”) repeats until either the maximum mismatch at the given iteration falls below a predecided threshold or the number of nodes crosses a fixed number of maximum nodes. In each iteration, the accuracy of the interpolants increases, which in turn enhances the meshfree approximation of the waveforms. Note that the choice of the ten worst points, as mentioned earlier, is completely arbitrary and chosen to accelerate the interpolant-generation process. To test the validity of the meshfree interpolated waveforms, we also performed a parameter estimation of a simulated BBH event using the meshfree interpolated waveforms and found posterior distributions to be consistent with the injected values. It is important to highlight that our method does not require importance sampling, as is used in deep-learning methods such as DINGO [49] and NESSAI [50].

A limitation we face involves the multiplication of SVD coefficients with the basis vectors. As the duration of the waveform increases, the computational expense of this multiplication rises, consequently impacting the evaluation cost of the meshfree interpolated waveform. In this context,

we can also use the empirical interpolation method (EIM) [38,51–53] to reduce the number of frequency points included in the basis vectors. It is not yet known whether this kind of strategy could be helpful within the meshfree framework in reducing the dimensionality of the basis vectors, which is especially needed for longer-duration waveforms. Further, in this work we used an iterative SVD scheme that updates the basis vectors with a new set of points by applying independent SVD of the updated amplitude and phase matrix. In the follow-up work, we will replace our iterative SVD scheme with advanced updated SVD algorithms [35,36] and compare the computation cost of performing independent SVD in each iteration and updating the initial basis and singular values from the first iteration using updated SVD algorithms. Another notable aspect is that a Euclidean distance measure is employed in Eq. (5b), even though the parameter space we operate in is inherently non-Euclidean. Nevertheless, we choose nodes in chirp-time coordinates ( $\theta_0$  and  $\theta_3$ ) instead of component masses or chirp mass and mass ratio as the metric varies slowly in chirp-time coordinates. Additionally, it is worth noting that this work solely focused on the aligned-spin waveform model NRHybSur3dq8. Subsequent extensions of this work will explore broadening the meshfree framework to incorporate NR surrogate models encompassing subdominant modes, considering spins, and addressing the aforementioned limitations. We also performed similar exercises for other aligned-spin waveform models, such as IMRPhenomXAS and SEOBNRv4\_ROM. Since these waveform models are already fast to evaluate, we do not get any significant speedup in their corresponding meshfree waveform evaluation. However, the speedup can be further enhanced by employing strategies such as adaptive frequency resolution [54] and the EIM procedure, as mentioned earlier. In future follow-ups of this work, we would also consider extending this meshfree framework to include NR surrogate models that include subdominant modes, precessing spins, and possibly eccentric models.

### ACKNOWLEDGMENTS

We thank Srashti Goyal for carefully going through the manuscript and giving helpful comments. We especially

thank the anonymous referee for their careful review and helpful suggestions. We thank Prayush Kumar for his help in the initial stages of this work. We also thank Abhishek Sharma and Sachin Shukla for their useful suggestions and comments. L. P. is supported by the Research Scholarship Program of Tata Consultancy Services (TCS). A. R. is supported by the research program of the Netherlands Organisation for Scientific Research (NWO). A. S. gratefully acknowledges the generous grant provided by the Department of Science and Technology, India, through the DST-ICPS (Interdisciplinary Cyber Physical Systems) cluster project funding. We thank the HPC support staff at IIT Gandhinagar for their help and cooperation. The authors are grateful for the computational resources provided by the LIGO Laboratory and supported by the National Science Foundation Grants No. PHY-0757058 and No. PHY-0823459. This material is based upon work supported by NSF’s LIGO Laboratory, which is a major facility fully funded by the National Science Foundation. This research has made use of data or software obtained from the Gravitational Wave Open Science Center [55], a service of the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA. This material is based upon work supported by NSF’s LIGO Laboratory, which is a major facility fully funded by the National Science Foundation, as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded through the European Gravitational Observatory (EGO), the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN), and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain. KAGRA. is supported by the Ministry of Education, Culture, Sports, Science and Technology (MEXT), Japan Society for the Promotion of Science (JSPS) in Japan; National Research Foundation (NRF) and Ministry of Science and ICT (MSIT) in Korea; Academia Sinica (AS) and National Science and Technology Council (NSTC) in Taiwan.

[1] 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).

[2] G. Pratten, S. Husa, C. García-Quirós, M. Colleoni, A. Ramos-Buades, H. Estellés, and R. Jaume, Setting the cornerstone for a family of models for gravitational waves from compact binaries: The dominant harmonic for nonprecessing quasicircular black holes, *Phys. Rev. D* **102**, 064001 (2020).



- [3] S. Khan, K. Chatziioannou, M. Hannam, and F. Ohme, Phenomenological model for the gravitational-wave signal from precessing binary black holes with two-spin effects, *Phys. Rev. D* **100**, 024059 (2019).
- [4] 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).
- [5] 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).
- [6] A. Bohé, L. Shao, A. Taracchini, A. Buonanno, S. Babak, I. W. Harry, I. Hinder, S. Ossokine, M. Pürrer, V. Raymond *et al.*, Improved effective-one-body model of spinning, nonprecessing binary black holes for the era of gravitational-wave astrophysics with advanced detectors, *Phys. Rev. D* **95**, 044028 (2017).
- [7] R. Cotesta, S. Marsat, and M. Pürrer, Frequency-domain reduced-order model of aligned-spin effective-one-body waveforms with higher-order modes, *Phys. Rev. D* **101**, 124040 (2020).
- [8] J. Blackman, S. E. Field, M. A. Scheel, C. R. Galley, C. D. Ott, M. Boyle, L. E. Kidder, H. P. Pfeiffer, and B. Szilágyi, Numerical relativity waveform surrogate model for generically precessing binary black hole mergers, *Phys. Rev. D* **96**, 024058 (2017).
- [9] 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).
- [10] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, L. E. Kidder, and H. P. Pfeiffer, Surrogate model of hybridized numerical relativity binary black hole waveforms, *Phys. Rev. D* **99**, 064045 (2019).
- [11] K. Cannon, C. Hanna, and D. Keppel, Interpolating compact binary waveforms using the singular value decomposition, *Phys. Rev. D* **85**, 081504 (2012).
- [12] P. Ajith, M. Hannam, S. Husa, Y. Chen, B. Brügmann, N. Dorband, D. Müller, F. Ohme, D. Pollney, C. Reisswig, L. Santamaría, and J. Seiler, Inspiral-merger-ringdown waveforms for black-hole binaries with nonprecessing spins, *Phys. Rev. Lett.* **106**, 241101 (2011).
- [13] A. J. K. Chua, C. R. Galley, and M. Vallisneri, Reduced-order modeling with artificial neurons for gravitational-wave inference, *Phys. Rev. Lett.* **122**, 211101 (2019).
- [14] K. G. Arun, A. Buonanno, G. Faye, and E. Ochsner, Higher-order spin effects in the amplitude and phase of gravitational waveforms emitted by inspiraling compact binaries: Ready-to-use gravitational waveforms, *Phys. Rev. D* **79**, 104023 (2009).
- [15] S. Khan and R. Green, Gravitational-wave surrogate models powered by artificial neural networks, *Phys. Rev. D* **103**, 064015 (2021).
- [16] D. Williams, I. Heng, J. Gair, J. Clark, and B. Khamesra, Precessing numerical relativity waveform surrogate model for binary black holes: A Gaussian process regression approach, *Phys. Rev. D* **101**, 063011 (2020).
- [17] J. Lee, S. H. Oh, K. Kim, G. Cho, J. J. Oh, E. J. Son, and H. M. Lee, Deep learning model on gravitational waveforms in merging and ringdown phases of binary black hole coalescences, *Phys. Rev. D* **103**, 123023 (2021).
- [18] R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari, V. B. Adya, C. Affeldt, D. Agarwal *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), Population of merging compact binaries inferred using gravitational waves through GWTC-3, *Phys. Rev. X* **13**, 011048 (2023).
- [19] 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).
- [20] 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).
- [21] E. T. Newman and R. Penrose, Note on the Bondi-Metzner-Sachs group, *J. Math. Phys. (N.Y.)* **7**, 863 (1966).
- [22] J. N. Goldberg, A. J. MacFarlane, E. T. Newman, F. Rohrlich, and E. C. G. Sudarshan, Spin  $s$  spherical harmonics and  $\tilde{O}$ , *J. Math. Phys. (N.Y.)* **8**, 2155 (1967).
- [23] 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).
- [24] 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).
- [25] 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).
- [26] 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).
- [27] 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).
- [28] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. J. Forteza, and A. Bohé, Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and anatomy of the signal, *Phys. Rev. D* **93**, 044006 (2016).
- [29] F. Kuo and J. Kaiser, *System Analysis by Digital Computer* (Wiley, New York, 1966).
- [30] M. Pürrer, Frequency-domain reduced order models for gravitational waves from aligned-spin compact binaries, *Classical Quantum Gravity* **31**, 195010 (2014).



- [31] 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, *Phys. Rev. D* **108**, 123040 (2023).
- [32] G. E. Fasshauer, *Meshfree Approximation Methods with Matlab* (World Scientific, Singapore, 2007), 10.1142/6437.
- [33] 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).
- [34] 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).
- [35] M. Brand, Fast low-rank modifications of the thin singular value decomposition, *Linear Algebra Appl.* **415**, 20 (2006), special Issue on Large Scale Linear and Nonlinear Eigenvalue Problems.
- [36] P. Stange, On the efficient update of the singular value decomposition, *Proc. Appl. Math. Mech.* **8** (2008).
- [37] 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).
- [38] 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).
- [39] 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).
- [40] 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).
- [41] 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.
- [42] A. J. Chua, C. R. Galley, and M. Vallisneri, Reduced-order modeling with artificial neurons for gravitational-wave inference, *Phys. Rev. Lett.* **122**, 211101 (2019).
- [43] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer Series in Statistics (Springer, New York, 2009).
- [44] T. Hines, Python package containing the tools necessary for radial basis function (RBF) applications (2015), <https://github.com/treverhines/RBF.git>.
- [45] A. Manzotti and A. Dietz, Prospects for early localization of gravitational-wave signals from compact binary coalescences with advanced detectors, [arXiv:1202.4031](https://arxiv.org/abs/1202.4031).
- [46] J. S. Speagle, DYNESTY: A dynamic nested sampling package for estimating Bayesian posteriors and evidences, *Mon. Not. R. Astron. Soc.* **493**, 3132 (2020).
- [47] E. Higson, W. Handley, M. Hobson, and A. Lasenby, Dynamic nested sampling: An improved algorithm for parameter estimation and evidence calculation, *Stat. Comput.* **29**, 891 (2019).
- [48] 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, 10.5281/zenodo.7600689 (2023).
- [49] M. Dax, S. R. Green, J. Gair, J. H. Macke, A. Buonanno, and B. Schölkopf, Real-time gravitational wave science with neural posterior estimation, *Phys. Rev. Lett.* **127**, 241103 (2021).
- [50] M. J. Williams, J. Veitch, and C. Messenger, Nested sampling with normalizing flows for gravitational-wave inference, *Phys. Rev. D* **103**, 103006 (2021).
- [51] M. Barrault, Y. Maday, N. C. Nguyen, and A. T. Patera, An “empirical interpolation” method: Application to efficient reduced-basis discretization of partial differential equations, *C. R. Math.* **339**, 667 (2004).
- [52] Y. Maday, N. C. Nguyen, A. T. Patera, and S. H. Pau, A general multipurpose interpolation procedure: The magic points, *Commun. Pure Appl. Anal.* **8**, 383 (2009).
- [53] S. Chaturantabut and D. C. Sorensen, Nonlinear model reduction via discrete empirical interpolation, *SIAM J. Sci. Comput.* **32**, 2737 (2010).
- [54] S. Morisaki, Accelerating parameter estimation of gravitational waves from compact binary coalescence using adaptive frequency resolutions, *Phys. Rev. D* **104**, 044062 (2021).
- [55] <https://gwosc.org/> (2024).