Surrogate model of hybridized numerical relativity binary black hole waveforms

Vijay Varma,[1,*] Scott E. Field,[2] Mark A. Scheel,[1] Jonathan Blackman,[1] Lawrence E. Kidder,[3] and Harald P. Pfeiffer[4]

[1]*Theoretical Astrophysics, California Institute of Technology, Pasadena, California 91125, USA*
[2]*Department of Mathematics, Center for Scientific Computing and Visualization Research, University of Massachusetts, Dartmouth, Massachusetts 02747, USA*
[3]*Center for Radiophysics and Space Research, Cornell University, Ithaca, New York 14853, USA*
[4]*Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Am Mühlenberg 1, Potsdam 14476, Germany*

(Received 28 January 2019; published 27 March 2019)

Numerical relativity (NR) simulations provide the most accurate binary black hole gravitational waveforms, but are prohibitively expensive for applications such as parameter estimation. Surrogate models of NR waveforms have been shown to be both fast and accurate. However, NR-based surrogate models are limited by the training waveforms' length, which is typically about 20 orbits before merger. We remedy this by hybridizing the NR waveforms using both post-Newtonian and effective one-body waveforms for the early inspiral. We present NRHybSur3dq8, a surrogate model for hybridized nonprecessing numerical relativity waveforms, that is valid for the entire LIGO band (starting at 20 Hz) for stellar mass binaries with total masses as low as 2.25 $M_\odot$. We include the $\ell \leq 4$ and (5, 5) spin-weighted spherical harmonic modes but not the (4, 1) or (4, 0) modes. This model has been trained against hybridized waveforms based on 104 NR waveforms with mass ratios $q \leq 8$, and $|\chi_{1z}|, |\chi_{2z}| \leq 0.8$, where χ_{1z} (χ_{2z}) is the spin of the heavier (lighter) black hole in the direction of orbital angular momentum. The surrogate reproduces the hybrid waveforms accurately, with mismatches $\lesssim 3 \times 10^{-4}$ over the mass range 2.25 $M_\odot \leq M \leq 300\ M_\odot$. At high masses ($M \gtrsim 40\ M_\odot$), where the merger and ringdown are more prominent, we show roughly 2 orders of magnitude improvement over existing waveform models. We also show that the surrogate works well even when extrapolated outside its training parameter space range, including at spins as large as 0.998. Finally, we show that this model accurately reproduces the spheroidal-spherical mode mixing present in the NR ringdown signal.

DOI: 10.1103/PhysRevD.99.064045

I. INTRODUCTION

The era of gravitational wave (GW) astronomy has been emphatically unveiled with the recent detections [1–7] by LIGO [8] and Virgo [9]. The detection of gravitational wave signals from compact binary sources is expected to become a routine occurrence as the advanced detectors reach their design sensitivity [10,11]. The possible science output from these events crucially depends on the availability of an accurate waveform model to compare against observed signals.

Numerical relativity (NR) is the only *ab initio* approach that accurately produces waveforms from the merger of a binary black hole (BBH) system. However, because NR simulations are computationally expensive, it is impractical to use them directly for applications such as parameter estimation, which can require upwards of 10^7 waveform evaluations. Therefore, the GW community has developed several approximate waveform models [12–21], some of which are fast to evaluate. These models make certain physically motivated assumptions about the underlying phenomenology of the waveforms, and they fit for any remaining free parameters using NR simulations.

Surrogate modeling [22,23] is an alternative approach that does not assume an underlying phenomenology and has been applied to a diverse range of problems [22–32]. NR surrogate models follow a data-driven approach, directly using the NR waveforms to implicitly reconstruct the underlying phenomenology. Three NR surrogate models have been built so far [26–28], including a seven-dimensional (mass ratio q and two spin vectors) model for generically precessing systems in quasicircular orbit [28]. Through cross-validation studies, these models were shown to be nearly as accurate as the NR waveforms they were trained against.

Despite the success of the surrogate modeling approach, existing surrogate models have two important limitations: (1) Because they are based solely on NR simulations, which typically are only able to cover the last ~20 orbits of a BBH inspiral, they are not long enough to span the full LIGO band

*vvarma@caltech.edu

for stellar mass binaries. (2) Apart from the first nonspinning model [26], these models have been restricted to mass ratios $q \leq 2$.[1] There are two reasons for this: (i) The 7D parameter space is vast, requiring at least a few thousand simulations to sufficiently cover it. (ii) Because of the smaller length scale introduced by the lighter black hole, NR simulations become increasingly more expensive with mass ratio.

In this work we address these limitations in the context of nonprecessing BBH systems. First, to include the early inspiral we "hybridize" the NR waveforms: each full waveform consists of a post-Newtonian (PN) and effective one-body (EOB) waveform at early times that is smoothly attached to a NR waveform at late times. Second, since we restrict ourselves to the three-dimensional space of nonprecessing BBHs, fewer simulations are necessary compared to the seven-dimensional case, and therefore we can direct computational resources to simulations with higher mass ratios. The resulting model, NRHybSur3dq8, is the first NR-based surrogate model to span the entire LIGO frequency band for stellar mass binaries; assuming a detector low-frequency cutoff of 20 Hz, this model is valid for total masses as low as 2.25 $M_\odot$. This model is based on 104 NR waveforms in the parameter range $q \leq 8$, and $|\chi_{1z}|, |\chi_{2z}| \leq 0.8$, where χ_{1z} (χ_{2z}) is the dimensionless spin of the heavier (lighter) black hole (BH).

The plus (h_+) and cross ($h_\times$) polarizations of GWs can be conveniently represented by a single complex time series, $\mathfrak{h} = h_+ - i h_\times$. The complex waveform on a sphere can be decomposed into a sum of spin-weighted spherical harmonic modes $\mathfrak{h}_{\ell m}$ [33,34], so that the waveform along any direction (ι, φ_0) in the binary's source frame is given by

$$\mathfrak{h}(t, \iota, \varphi_0) = \sum_{\ell=2}^{\infty} \sum_{m=-l}^{l} \mathfrak{h}_{\ell m}(t) {}^{-2}Y_{\ell m}(\iota, \varphi_0), \tag{1}$$

where ${}^{-2}Y_{\ell m}$ are the spin $= -2$ weighted spherical harmonics, ι is the inclination angle between the orbital angular momentum of the binary and line of sight to the detector, and φ_0 is the initial binary phase. φ_0 can also be thought of as the azimuthal angle between the x axis of the source frame and the line of sight to the detector. We define the source frame as follows: The z axis is along the orbital angular momentum direction, which is constant for nonprecessing BBH. The x axis is along the line of separation from the lighter BH to the heavier BH at some reference time or frequency. The y axis completes the triad.

The $\ell = |m| = 2$ terms typically dominate the sum in Eq. (1) and are referred to as the "'quadrupole" modes. Studies [35–42] have shown that the nonquadrupole modes, while being subdominant, can play a non-negligible role in detection and parameter estimation of GW sources, particularly for large signal to noise ratio (SNR), large total mass, large mass ratio, or large inclination angle ι. For the first event, GW150914 [1], the systematic errors due to the quadrupole-mode-only approximation are generally smaller than the statistical errors [43,44], although higher modes may lead to modest changes in some of the extrinsic parameter values [45]. However, as the detectors approach their design sensitivity [10], one should prepare for high-SNR sources (particularly at larger mass ratios than those seen so far), where the quadrupole-mode-only approximation breaks down. In addition, nonquadrupole modes can help break the degeneracy between the binary inclination and distance, which is present for quadrupole-mode-only models (see, e.g., [14,46,47]).

In this work, we model the following spin-weighted spherical harmonic modes: $\ell \leq 4$ and (5, 5), but not the (4, 1) or (4, 0) modes.[2] Several inspiral-merger-ringdown waveform models [14,15,20,21] that include nonquadrupole modes have been developed in recent years; however, compared to those models we show an improved accuracy and we include more modes.

The rest of the paper is organized as follows. In Sec. II, we choose the parameters at which to perform NR simulations, which will be used for training the surrogate model. Section III describes the NR simulations. Section IV describes our procedure to compute the waveform for the early inspiral using PN and EOB waveforms. Section V describes our hybridization procedure to attach the early inspiral waveform to the NR waveforms. Section VI describes the construction of the surrogate model. In Sec. VII, we test the surrogate model by comparing against NR and hybrid waveforms. We end with some concluding remarks in Sec. VIII. We make our model available publicly through the easy-to-use PYTHON package *gwsurrogate* [48]. In addition, our model is implemented in C with PYTHON wrapping in the LIGO algorithm library [49]. We provide an example PYTHON evaluation code in [50].

II. TRAINING SET GENERATION

A. Greedy parameters from PN surrogate model

We do not know *a priori* the distribution or number of NR simulations required to build an accurate surrogate model. Furthermore, we hope to select a representative distribution that will allow for an accurate surrogate to be built with as few NR simulations as possible. Therefore, we estimate this distribution by first building a surrogate model for PN waveforms; we find that parameters suitable for building an accurate PN surrogate are also suitable for building a NR or a hybrid NR-PN surrogate.

[1]We use the convention $q = m_1/m_2$, where m_1 and m_2 are the masses of the component black holes, with $m_1 \geq m_2$.

[2]Because of the symmetries of nonprecessing BBHs [see Eq. (23)], the $m < 0$ modes contain the same information as the $m > 0$ modes and do not need to be modeled separately.

We use the same methods to build the PN surrogate as we use for the hybrid surrogate (cf. Sec. VI). We use the PN waveforms described in Sec. IV A; however, for simplicity we only model the (2, 2) mode. In addition, we restrict the length of the PN waveforms to be $5000M$, terminating at the innermost-stable-circular-orbit's orbital frequency, $\omega_{\rm orb} = 6^{-3/2}$ rad$/M$, where M is the total mass of the binary.

We determine the desired training dataset of parameters as follows. We begin with just the corner cases of the parameter space; for the 3D case considered here, that consists of eight points at $(q, \chi_{1z}, \chi_{2z}) = (1 \text{ or } 8, \pm 0.8, \pm 0.8)$. We build up the desired set of parameters iteratively, in a greedy manner: At each iteration we build a PN surrogate using the current training dataset and test the model against a much larger ($\sim$10 times) validation dataset. The validation dataset is generated by randomly resampling the parameter space at each iteration. Since the boundary cases are expected to be more important, for 30% of the points in the validation set we sample only from the boundary of the parameter space, which corresponds to the faces of a cube in the 3D case. We select the parameter in the validation set that has the largest error [cf. Eq. (2)] and add this to our training set (hence the name greedy parameters). We repeat until the validation error reaches a certain threshold.

In order to estimate the difference between two complexified waveforms, $\mathfrak{h}_1$ and $\mathfrak{h}_2$, we use the time-domain mismatch

$$\mathcal{MM} = 1 - \frac{\langle \mathfrak{h}_1, \mathfrak{h}_2 \rangle}{\sqrt{\langle \mathfrak{h}_1, \mathfrak{h}_1 \rangle \langle \mathfrak{h}_2, \mathfrak{h}_2 \rangle}}, \tag{2}$$

$$\langle \mathfrak{h}_1, \mathfrak{h}_2 \rangle = \left| \int_{t_{\rm min}}^{t_{\rm max}} \mathfrak{h}_1(t) \mathfrak{h}_2^*(t) dt \right|, \tag{3}$$

where $*$ indicates a complex conjugation, and $|.|$ indicates the absolute value. Note that in this section, we do not perform an optimization over time and phase shifts. In addition, we assume a flat noise curve.

Figure 1 shows how the maximum validation error decreases as we add greedy parameters to our training dataset. For our case, we stop at 100 greedy parameters (at which point the mismatch is $< 10^{-6}$) and use those parameters to perform the NR simulations. Note that we do not expect 100 NR simulations to produce a NR surrogate with comparable accuracy, $\mathcal{MM} < 10^{-6}$, for two reasons. First, unlike the PN waveforms used here, the NR simulations also include the merger-ringdown part, which we expect to be more difficult to model. Second, the NR numerical truncation error is typically higher than 10^{-6} in mismatch, therefore the numerical noise will limit the accuracy.

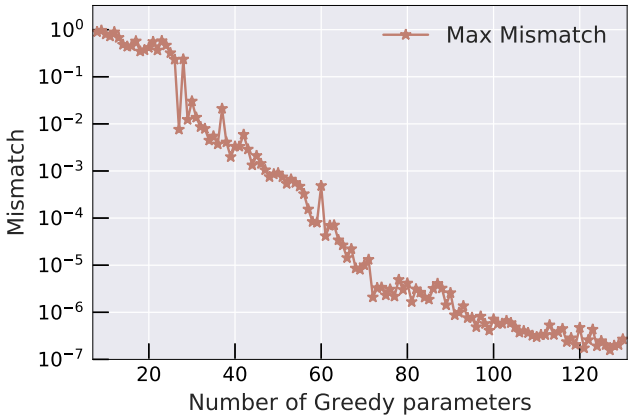

FIG. 1. Largest mismatch of the surrogate (over the entire validation set) as a function of number of greedy parameters used to train the PN surrogate. The PN surrogate is seen to converge to the validation waveforms as the size of the training dataset increases.

III. NR SIMULATIONS

The NR simulations for this model are performed using the Spectral Einstein Code (SpEC) [51–56] developed by the SXS [57] Collaboration. Of the 100 cases determined in Sec. II, only 91 simulations were successfully completed.[3] These simulations have been assigned the identifiers SXS:BBH:1419–SXS:BBH:1509 and are made publicly available through the SXS public catalog [58]. For cases with equal mass, but unequal spins, we can exchange the two BHs to get an extra data point. There are 13 such cases, leading to a total of 104 NR waveforms. These are shown as circular markers in Fig. 2.

The start time of these simulations varies between $4270M$ and $5227M$ before the peak of the waveform amplitude [defined in Eq. (38)], where $M = m_1 + m_2$ is the total Christodoulou mass measured after the initial burst of junk radiation. The algorithm for choosing a fiducial time at which junk radiation ends is discussed in Ref. [59]. The initial orbital parameters are chosen through an iterative procedure [60] such that the orbits are quasicircular; the largest eccentricity for these simulations is 7.5×10^{-4}, while the median value is 4.2×10^{-4}. The waveforms are extracted at several extraction surfaces at varying finite radii from the origin and then extrapolated to future-null infinity [61]. Finally, the extrapolated waveforms are corrected to account for the initial drift of the center of mass [62,63]. The time steps during the simulations are chosen nonuniformly using an adaptive time stepper [59]. We interpolate these data to a uniform time step of $0.1M$; this is dense enough to capture all frequencies of interest, including near merger.

[3]The main reason for failure is large constraint violation as the binary approaches merger. We believe a better gauge condition may be needed for some of these simulations.

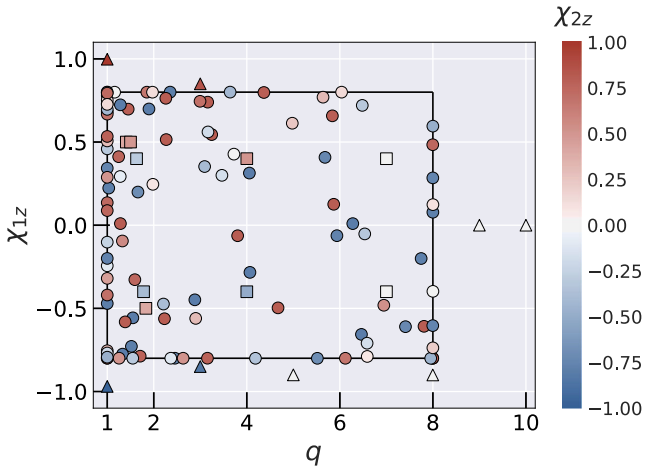

FIG. 2. The parameter space covered by the 104 NR waveforms (circle markers) used in the construction of the surrogate model in Sec. VI. We also show the nine long NR waveforms (square markers) used to test hybridization in Sec. VII B and the eight NR waveforms (triangle markers) used to test extrapolation in Sec. VII C. The axes show the mass ratio and the spin on the heavier BH, while the colors indicate the spin on the lighter BH. The black rectangle indicates the bounds of the training region: $1 \leq q \leq 8, -0.8 \leq \chi_{1z}, \chi_{2z} \leq 0.8$.

IV. EARLY INSPIRAL WAVEFORMS

While NR provides accurate waveforms, computational constraints limit NR to only the late inspiral, merger, and ringdown phases. Fortunately, PN/EOB waveforms are expected to be accurate in the early inspiral. Hence we can "stitch" together an early inspiral waveform and a NR waveform, to get a *hybrid* waveform [40,64–71] that spans the entire frequency range relevant for ground-based detectors. In this section, we describe the waveforms we use for the early inspiral, leaving the hybridization procedure for the next section.

A. PN waveforms

We first generate PN waveforms as implemented in the GWFRAMES package [72]. For the orbital phase we include nonspinning terms up to 4 PN order [73–77] and spin terms up to 2.5 PN order [78–80]. We use the TaylorT4 [81] approximant to generate the PN phase; however, as described below, we replace this phase with an EOB-derived phase. For the amplitudes, we include terms up to 3.5 PN order [82–84].

The spherical harmonic modes of the PN waveform can be written (after rescaling to unit total mass and unit distance) as [74,82]

$$\mathfrak{h}^{\rm PN}_{\ell m} = 2\eta (v^{\rm PN})^2 \sqrt{\frac{16\pi}{5}} H^{\rm PN}_{\ell m} e^{-im\phi^{\rm PN}_{\rm orb}}, \tag{4}$$

where $\eta = q/(1+q)^2$ is the symmetric mass ratio, $v^{\rm PN}$ is the characteristic speed that sets the perturbation scale in PN, $\phi^{\rm PN}_{\rm orb}$ is the (real) orbital phase, and $H^{\rm PN}_{\ell m}$ are the complex amplitudes of different modes. Note that we ignore the tail distortions [85,86] to the orbital phase as these are 4 PN corrections (see, e.g., [87]).

The complex strain $\mathfrak{h}^{\rm PN}_{\ell m}$ is obtained as a time series from GWFRAMES. We can absorb the complex part of the amplitudes into the phases and rewrite the strain as

$$\mathfrak{h}^{\rm PN}_{\ell m} = A^{\rm PN}_{\ell m} e^{-i\phi^{\rm PN}_{\ell m}}, \tag{5}$$

$$\phi^{\rm PN}_{\ell m} = m\phi^{\rm PN}_{\rm orb} + \xi^{\rm PN}_{\ell m}, \tag{6}$$

$$\phi^{\rm PN}_{\rm orb} = \frac{\phi^{\rm PN}_{22}}{2}, \tag{7}$$

where $A^{\rm PN}_{\ell m}$ and $\phi^{\rm PN}_{\ell m}$ are the real amplitude and phase of a given mode, and $\xi^{\rm PN}_{\ell m}$ is an offset that captures the complex part of $H^{\rm PN}_{\ell m}$. Note that Eqs. (6) and (7) together imply $\xi^{\rm PN}_{22} = 0$; $H^{\rm PN}_{22}$ contains complex terms starting at 2.5 PN, but these appear as 5 PN corrections in the phase (see, e.g., [87]), which we can safely ignore.

At this stage, $A^{\rm PN}_{\ell m}$, $\phi^{\rm PN}_{\ell m}$, and $\xi^{\rm PN}_{\ell m}$ are functions of time. But they can be recast as functions of the characteristic speed by first computing

$$v^{\rm PN}(t) = \left(\frac{d\phi^{\rm PN}_{\rm orb}}{dt}\right)^{1/3}, \tag{8}$$

where the derivative is performed numerically, and then inverting Eq. (8) to obtain $t(v^{\rm PN})$. Then we define

$$A^{\rm PN}_{\ell m}(v^{\rm PN}) = |\mathfrak{h}^{\rm PN}_{\ell m}(t(v^{\rm PN}))|, \tag{9}$$

$$\xi^{\rm PN}_{\ell m}(v^{\rm PN}) = \phi^{\rm PN}_{\ell m}(t(v^{\rm PN})) - m\phi^{\rm PN}_{\rm orb}(t(v^{\rm PN})). \tag{10}$$

Note that the PN waveform is generated in the source frame defined such that the reference time is the initial time. This also ensures that the heavier BH is on the positive x axis at the initial time, and the initial orbital phase is zero.

To summarize, from the GWFRAMES package, we obtain the complex time series $\mathfrak{h}^{\rm PN}_{\ell m}$ [Eq. (5)]. We compute the orbital phase [Eq. (7)], the real amplitudes [Eq. (9)], and the phase offsets [Eq. (10)]. These three quantities are obtained as a time series but can be represented as functions of the characteristic speed using Eq. (8).

B. EOB correction

As was shown in previous works [35,40], we find that the accuracy of the inspiral waveform can be improved by replacing the PN phase with the phase derived from a NR-calibrated EOB model. For this work we use SEOBNRv4 [17].

SEOBNRv4 is a time-domain model that includes only the (2, 2) mode, which we can decompose as follows:

$$\mathfrak{h}_{22}^{\rm EOB} = A_{22}^{\rm EOB} e^{-i\phi_{22}^{\rm EOB}}, \tag{11}$$

where $A_{22}^{\rm EOB}$ and $\phi_{22}^{\rm EOB}$ are the real amplitude and phase of the (2, 2) mode. These are functions of time, but following the same procedure as earlier, they can be recast in terms of the characteristic speed,

$$\phi_{\rm orb}^{\rm EOB}(t) = \frac{\phi_{22}^{\rm EOB}(t)}{2}, \tag{12}$$

$$v^{\rm EOB}(t) = \left(\frac{d\phi_{\rm orb}^{\rm EOB}}{dt}\right)^{1/3}, \tag{13}$$

where the derivative is performed numerically, and we invert Eq. (13) to obtain $t(v^{\rm EOB})$. We replace $v^{\rm PN} \to v^{\rm EOB}$ in Eqs. (9) and (10) to get, respectively, the EOB-corrected amplitudes and phase offsets

$$A_{\ell m}^{\rm ins}(t) = A_{\ell m}^{\rm PN}(v^{\rm EOB}(t)), \tag{14}$$

$$\xi_{\ell m}^{\rm ins}(t) = \xi_{\ell m}^{\rm PN}(v^{\rm EOB}(t)). \tag{15}$$

Note that, in practice, computing $A_{\ell m}^{\rm ins}(t)$ and $\xi_{\ell m}^{\rm ins}(t)$ is accomplished via an interpolation in v: $A_{\ell m}^{\rm PN}(v)$ and $\xi_{\ell m}^{\rm PN}(v)$ as computed in Eqs. (9) and (10) are known only at particular values of v, which are $v^{\rm PN}(t_{i_{\rm PN}})$, where $t_{i_{\rm PN}}$ are the times in the PN time series; we interpolate $A_{\ell m}^{\rm PN}(v)$ and $\xi_{\ell m}^{\rm PN}(v)$ to the points $v^{\rm EOB}(t_{i_{\rm EOB}})$, where $t_{i_{\rm EOB}}$ are the times in the EOB time series. We use a cubic-spline interpolation scheme as implemented in SciPy [88].

Following Eq. (6), the EOB-corrected phases are given by

$$\phi_{\ell m}^{\rm ins} = m\phi_{\rm orb}^{\rm EOB} + \xi_{\ell m}^{\rm ins}, \tag{16}$$

where we use the EOB orbital phase from Eq. (12). Finally, our EOB-corrected inspiral waveform modes are given by

$$\mathfrak{h}_{\ell m}^{\rm ins} = A_{\ell m}^{\rm ins} e^{-i\phi_{\ell m}^{\rm ins}}. \tag{17}$$

Figure 3 shows an example of PN and EOB-corrected waveforms along with the corresponding NR waveform. All three waveforms have the same starting orbital frequency and their initial orbital phase is set to zero. We see that the PN waveform becomes inaccurate at late times, as expected. The EOB-corrected waveform, on the other hand, remains faithful to the NR waveform until much later times.

V. HYBRIDIZATION

In this section we describe our procedure to stitch together an inspiral waveform (described in Sec. IV) to a NR waveform (described in Sec. III).

We start by generating inspiral and NR waveforms with the same component masses and spins. We note that the spins measured in SpEC simulations agree well with PN theory [89]. However, the PN and NR waveforms are typically represented in different coordinate systems that need to be aligned with each other as follows. The two coordinate systems are related to each other by a possible time translation and a possible rotation by three Euler

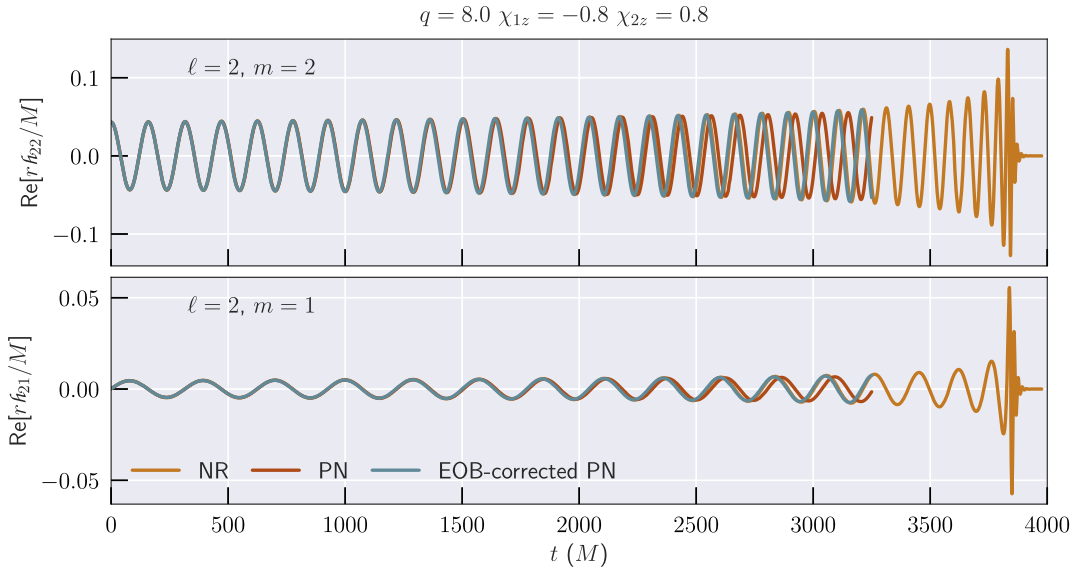

FIG. 3. NR, PN (Sec. IV A), and EOB-corrected PN (Sec. IV B) waveforms for an example case. We show the (2, 2) and (2, 1) modes. The binary parameters are shown at the top of the plot. The EOB-corrected PN waveform [35,40] stays faithful to the NR waveform until much later times, compared to the pure PN waveform.

angles: inclination angle ι, initial binary phase φ_0, and polarization angle ψ. For nonprecessing BBH, the first angle ι is trivially specified by requiring that the z axis is along the direction of orbital angular momentum. This leaves us with the freedom to vary φ_0 and ψ. We choose the hybridization frame and time shifts by minimizing a cost function in a suitable matching region; this is described in more detail below.

A. Choice of cost function

We use the following cost function when comparing two waveforms, $\mathfrak{h}$ and $\tilde{\mathfrak{h}}$, in the matching region:

$$\mathcal{E}[\mathfrak{h},\tilde{\mathfrak{h}}] = \frac{1}{2}\frac{\sum_{\ell,m}\int_{t_1}^{t_2}|\mathfrak{h}_{\ell m}(t)-\tilde{\mathfrak{h}}_{\ell m}(t)|^2 dt}{\sum_{\ell,m}\int_{t_1}^{t_2}|\mathfrak{h}_{lm}(t)|^2 dt}, \tag{18}$$

where t_1 and t_2 denote the start and end of the matching region, to be defined in Sec. V C, and the sum does not include $m=0$ modes for reasons described in Sec. V B. This cost function was introduced in Ref. [27] and is shown to be related to the weighted average of the mismatch over the sky.

We minimize the cost function by varying the time and frame shifts between the NR and inspiral waveforms

$$\min_{t_0,\varphi_0,\psi}\mathcal{E}[\mathfrak{h}^{\rm NR}(t;\varphi_0,\psi),\mathfrak{h}^{\rm ins}(t;t_0)], \tag{19}$$

$$\mathfrak{h}^{\rm NR}_{\ell m}(t;\varphi_0,\psi) = \mathfrak{h}^{\rm NR}_{\ell m}(t)e^{im\varphi_0}e^{2i\psi}, \tag{20}$$

$$\mathfrak{h}^{\rm ins}_{\ell m}(t;t_0) = \mathfrak{h}^{\rm ins}_{\ell m}(t-t_0). \tag{21}$$

We perform the time shifts on the inspiral waveform so that the matching region always corresponds to the same segment of the NR waveform. The frame shifts are performed on the NR waveform so as to preserve the initial frame alignment of the inspiral waveform (cf. Sec. IV A). This alignment gets inherited by the hybrid waveform and is important in the surrogate construction.

B. $m=0$ modes

We find that the $m=0$ modes of the inspiral waveforms do not agree very well with the NR waveforms. There are several possible reasons for this [90]: (1) The NR waveform does not have the correct "memory" contribution since this depends on the entire history of the system starting at $t=-\infty$, while the NR simulation covers only the last few orbits. (2) The extrapolation to future-null infinity does not work as well for these modes [59]. This could be improved in the future with Cauchy characteristic extraction (CCE) [91–94]. (3) The amplitude of these modes is very small except very close to merger; therefore the early part of the NR waveform where we compare with the inspiral waveforms is contaminated by numerical noise.

Therefore, when constructing the hybrid waveforms, we set the entire inspiral waveform to zero for these modes,

$$\mathfrak{h}^{\rm ins}_{\ell,m=0} = 0. \tag{22}$$

When computing the cost function [Eq. (18)], we ignore the $m=0$ modes.

This means that our hybrid waveforms for these modes are equivalent to the NR waveforms. In addition, the main contribution for these modes comes from the region close to merger, which does not correspond to a memory signal, but instead is due to axisymmetric excitations near merger (cf. bottom panel of Fig. 5).

C. Choice of matching region

There are several considerations to take into account when choosing a matching region $[t_1,t_2]$ for the cost function [Eq. (18)]: (1) The NR and inspiral waveforms should agree with each other reasonably in this region; at early times the NR waveform is contaminated by junk radiation, while at late times the inspiral waveform deviates from NR (cf. Figs. 3 and 5). (2) The matching region should be wide enough that the cost function is meaningful.

Our matching region starts at $1000M$ after the start of the NR waveform; we find that this is necessary to avoid noise due to junk radiation in some of the higher-order modes. The length of the NR waveforms from the start of the matching region to the peak of the waveform amplitude varies between $3270M$ and $4227M$. The width of the matching region is then chosen to be equal to the time taken for 3 orbits of the binary. We use the phase of the (2, 2) mode of the NR waveform to determine this. This choice ensures the width of the matching region scales appropriately with the NR starting frequency, so that we get wider matching regions when the NR waveform starts early in the inspiral.

D. Allowed ranges for frame and time shifts

The allowed range for φ_0 is $[0,2\pi]$. For nonprecessing binaries the allowed values for ψ can be restricted by taking into account the symmetries of the system. We will show that this restriction is a consequence of the well-known relationship

$$\mathfrak{h}_{\ell,-m} = (-1)^{\ell}\mathfrak{h}^{*}_{\ell,m}, \tag{23}$$

between the $m<0$ modes and the $m>0$ modes for nonprecessing binaries orbiting in the $x-y$ plane [95]. We compute the shifted waveform

$$\begin{aligned}\mathfrak{h}_{\ell,-m}(t)e^{-im\varphi_0}e^{2i\psi} &= \mathfrak{h}_{\ell,-m}(t;\varphi_0,\psi)\\ &= (-1)^{\ell}(\mathfrak{h}_{\ell,m}(t;\varphi_0,\psi))^*\\ &= (-1)^{\ell}e^{-2i\psi}e^{-im\varphi_0}\mathfrak{h}^*_{\ell,m}(t)\\ &= e^{-2i\psi}e^{-im\varphi_0}\mathfrak{h}_{\ell,-m}(t)\\ \Rightarrow e^{2i\psi} &= e^{-2i\psi}. \end{aligned} \tag{24}$$

Equation (24) implies that the only allowed values for ψ are 0 and $\pi/2$.[4] If the inspiral waveform and the NR waveform have the same sign convention, then $\psi = 0$. Unfortunately, not all NR catalogs and PN-waveform codes use the same sign convention, so we allow the possibility of $\psi = \pi/2$ to account for this.

To set the allowed range for t_0, we begin by computing the orbital frequency of the inspiral waveform $\omega^{\rm ins}$ as half the frequency of the (2, 2) mode. Similarly, we compute the orbital frequency of the NR waveform $\omega^{\rm NR}$. We first time align the NR and inspiral waveforms such that their frequencies match at the start of the matching region. This gives us a good starting point to vary the time shift.

We also define

$$\omega^{\rm ins}_{\rm mid} = \omega^{\rm NR}(t = t_1), \tag{25}$$

$$\omega^{\rm ins}_{\rm low} = 0.995 \times \omega^{\rm ins}_{\rm mid}, \tag{26}$$

$$\omega^{\rm ins}_{\rm hi} = 1.005 \times \omega^{\rm ins}_{\rm mid}, \tag{27}$$

where $\omega^{\rm NR}(t = t_1)$ is the NR frequency at the start of the matching region. The allowed range for time shifts t_0 is restricted to lie in the interval $[t^{\rm ins}_{\rm low} - t^{\rm ins}_{\rm mid}, t^{\rm ins}_{\rm hi} - t^{\rm ins}_{\rm mid}]$, where $t^{\rm ins}_{\rm low}$, $t^{\rm ins}_{\rm mid}$, and $t^{\rm ins}_{\rm hi}$ are the times at which $\omega^{\rm ins}(t)$ is equal to $\omega^{\rm ins}_{\rm low}$, $\omega^{\rm ins}_{\rm mid}$, and $\omega^{\rm ins}_{\rm hi}$, respectively. In other words, the allowed range for t_0 is a region near $t_0 = 0$. $t_0 = 0$ is the case when the frequencies of the inspiral and the NR waveforms match at t_1, the start of the matching region. The lower (upper) limit for t_0 is chosen such that the inspiral waveform has a frequency equal to 0.995 (1.005) times the NR frequency at t_1.

The factors in Eqs. (25) and (27) are chosen such that the time shift that minimizes the cost function is always well within the range of allowed time shifts. Hence, choosing a wider range (i.e., values of these factors farther from unity) does not improve the hybridization procedure. Note also that, like the width of the matching region in Sec. V C, setting the range of time shifts based on the orbital frequency ensures that it scales appropriately with the start frequency of the NR waveform.

The minimization in Eq. (19) is performed as follows. We vary the time shift t_0 over 500 uniformly spaced values in the above mentioned time range.[5] For each of these time shifts t_0, we try both allowed values of $\psi \in \{0, \pi/2\}$. For each t_0 and ψ, we minimize the cost function over φ_0 using the Nelder-Mead downhill simplex minimization algorithm as implemented in SciPy [88]. To avoid local minima in the φ_0 minimization, we perform ten searches with different initial guesses, which are sampled from a uniform random distribution in the range $[0, 2\pi]$.

E. Stitching NR and inspiral waveforms

Having obtained the right frame and time shifts between the NR and inspiral waveforms, the final step is to smoothly stitch the inspiral waveform to the shifted NR waveform. The stitching is done using a smooth blending function

$$\tau(t) = \begin{cases} 0, & \text{if } t < t_1 \\ \sin^2\left(\frac{\pi}{2}\frac{t-t_1}{t_2-t_1}\right), & \text{if } t_1 \le t \le t_2 \\ 1, & \text{if } t > t_2, \end{cases} \tag{28}$$

where t_1 and t_2 take on the same values as those appearing in Eq. (18). Different blending functions have been proposed in the literature [64,67,69,96]. Our choice is equivalent to the blending function defined in Ref. [67]. We find that our results are not sensitive to the choice of blending function.

In what follows, for brevity, we drop the hybridization parameters φ_0, ψ, t_0 with the understanding that the models are stitched together after transforming into the hybridization frame

$$\mathfrak{h}^{\rm ins}_{\ell m}(t) \equiv \mathfrak{h}^{\rm ins}_{\ell m}(t; t_0), \tag{29}$$

$$\mathfrak{h}^{\rm NR}_{\ell m}(t) \equiv \mathfrak{h}^{\rm NR}_{\ell m}(t; \varphi_0, \psi). \tag{30}$$

Given the shifted waveforms and the blending function, there are still several ways in which one can stitch the waveforms together.

1. Inertial frame stitching

One could work with the complex waveform strain and define

$$\mathfrak{h}^{\rm Hyb}_{\ell m} = (1 - \tau(t))\mathfrak{h}^{\rm ins}_{\ell m}(t) + \tau(t)\mathfrak{h}^{\rm NR}_{\ell m}(t). \tag{31}$$

With this choice, by construction, the complex strain transitions smoothly from the inspiral part to the NR part over the matching region. However, the transition is more complicated for the frequency, since it involves time derivatives of the complex argument of the strain; the time

[4] $\psi = \pi$ is also allowed, but it is degenerate with $\psi = 0$.

[5] We find that increasing the number of time samples results in no noticeable improvement; the typical values of the cost function after minimization with 500 samples are $\mathcal{E} \sim 10^{-5}$, and using 1000 samples results in changes only of order $\Delta\mathcal{E} \lesssim 10^{-8}$.

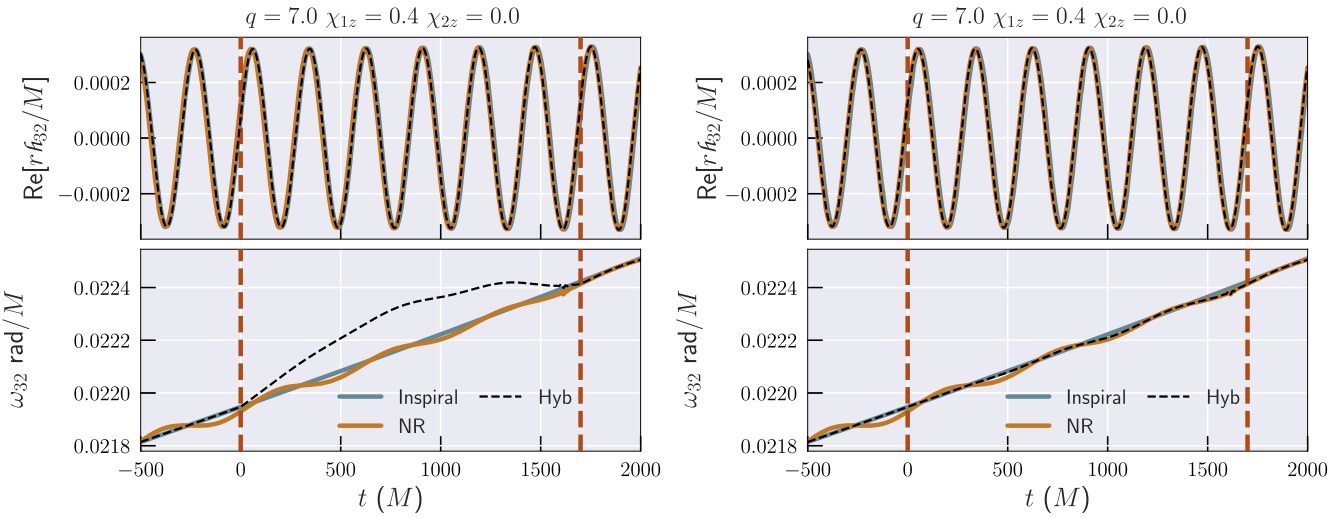

FIG. 4. (Left) The real part (top) and frequency (bottom) of the (3, 2) mode using the inertial frame stitching described in Sec. V E 1. The binary parameters are shown on the top of the plot. The vertical red dashed lines indicate the matching region. Note that this plot shows the inspiral and NR waveforms after the time and frame shifts are performed. (Right) Same, but using the amplitude-frequency stitching described in Sec. V E 2. Now we see that the frequency of the hybrid waveform agrees much better with the NR and inspiral data.

derivatives of the blending function do not behave like a smooth blending function. This is demonstrated in the left panel of Fig. 4: the inspiral and NR frequencies agree well in the matching region but the frequency of the hybrid waveform deviates from this.

2. Amplitude-frequency stitching

To avoid the undesirable artifacts described above, we choose to perform the inspiral-NR stitching using the amplitude and frequency rather than the inertial frame strain.

We begin by decomposing the NR and inspiral waveforms into their respective amplitude and phase,

$$\mathfrak{h}^{\rm NR}_{\ell m}(t) = A^{\rm NR}_{\ell m} e^{-i\phi^{\rm NR}_{\ell m}}, \qquad \mathfrak{h}^{\rm NR}_{\ell m}(t) = A^{\rm ins}_{\ell m} e^{-i\phi^{\rm ins}_{\ell m}}. \tag{32}$$

The frequency of each mode

$$\omega^{\rm NR}_{\ell m} = \frac{d\phi^{\rm NR}_{\ell m}}{dt}, \qquad \omega^{\rm ins}_{\ell m} = \frac{d\phi^{\rm ins}_{\ell m}}{dt}, \tag{33}$$

is then numerically computed from fourth-order finite difference approximations to the time derivative. Finally, we stitch the amplitude and frequency of each mode to get their hybrid versions,

$$A^{\rm Hyb}_{\ell m} = (1-\tau(t)) \quad A^{\rm ins}_{\ell m}(t) + \tau(t) A^{\rm NR}_{\ell m}, \tag{34}$$

$$\omega^{\rm Hyb}_{\ell m} = (1-\tau(t)) \quad \omega^{\rm ins}_{\ell m}(t) + \tau(t) \omega^{\rm NR}_{\ell m}. \tag{35}$$

To get the inertial frame strain we first need to integrate the frequency to get the phase. However, we already know the phase in the region before (only inspiral) and after (only NR) the matching region. So, we integrate the hybrid frequency

$$\phi^{\text{Hyb-match-region}}_{\ell m} = \int_{t_1}^{t_2} \omega^{\rm Hyb}_{\ell m} dt, \tag{36}$$

in the matching region using a fourth-order accurate Runge-Kutta scheme.

Finally, we set the phase of the hybrid waveform to

$$\phi^{\rm Hyb}_{\ell m} = \begin{cases} \phi^{\rm ins}_{\ell m} + \delta^1_{\ell m}, & \text{if } t < t_1 \\ \phi^{\text{Hyb-match-region}}_{\ell m} + \delta^2_{\ell m}, & \text{if } t_1 \le t \le t_2 \\ \phi^{\rm NR}_{\ell m}, & \text{if } t > t_2, \end{cases} \tag{37}$$

where $\delta^1_{\ell m}$ and $\delta^2_{\ell m}$ are chosen such that $\phi^{\rm Hyb}_{\ell m}$ is continuous at t_1 and t_2.

Since, by construction, the frequency transitions smoothly from the inspiral waveform to NR data, we eliminate the artifact seen in the bottom left panel of Fig. 4 (dashed line), as demonstrated in the right panel of Fig. 4.

We note that since the $m = 0$ modes are purely real (imaginary) and nonoscillatory for nonprecessing systems, they do not have a frequency associated with them, therefore we use the inertial frame stitching of Sec. V E 1 for these modes. For these modes the waveform goes from zero to the NR value over the matching region.

In the hybridized waveform we include the $\ell \le 4$ and (5, 5) modes, but not the (4, 1) or (4, 0) modes. For the (4, 1) and (4, 0) modes we find that the inspiral and NR waveforms do not agree very well. This is possibly due to issues in the extrapolation to future-null infinity [61] for these modes and could be resolved in the future with CCE

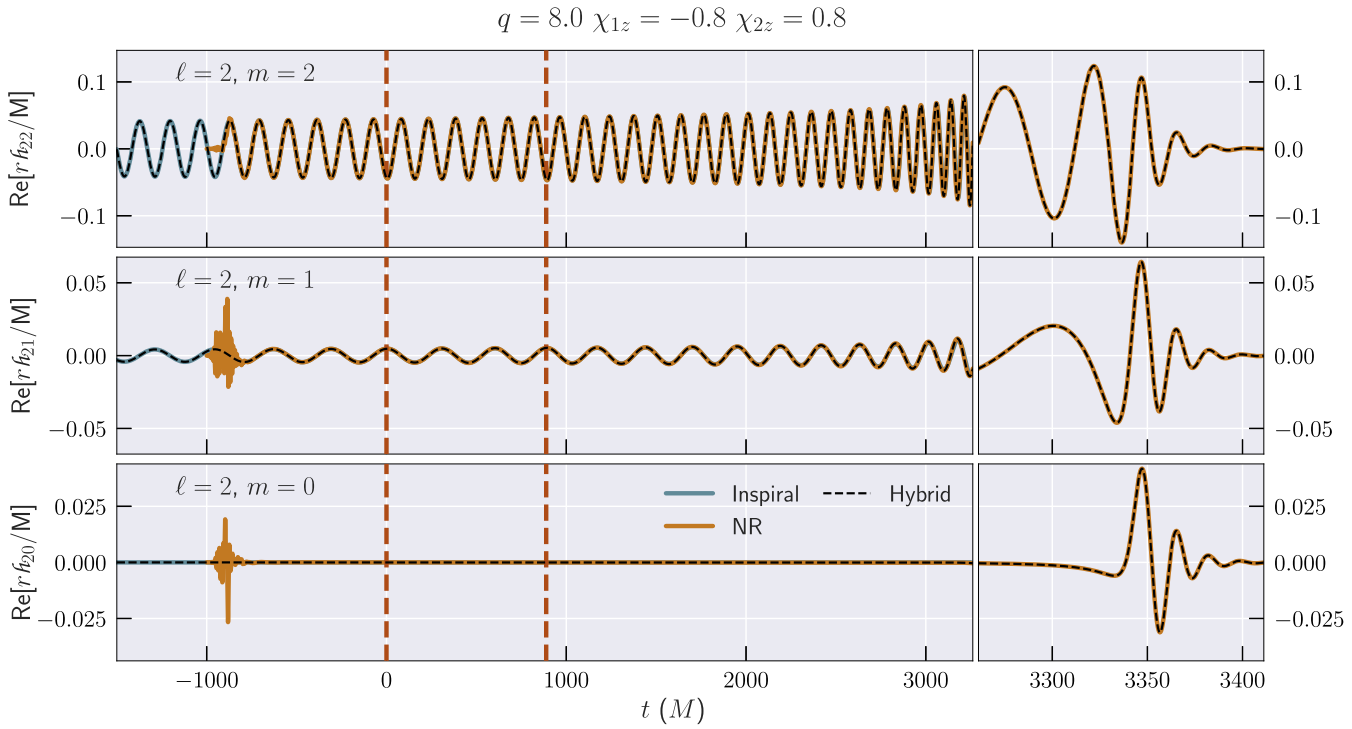

FIG. 5. An example hybrid waveform used in this work. We show the $\ell = 2$ modes of the inspiral, NR, and hybrid waveforms. The binary parameters are shown on the top of the plot. The vertical red dashed lines indicate the matching region. Note that this plot shows the inspiral and NR waveforms after the time and frame shifts are done.

[91–94] An example of the final NR, inspiral, and hybrid waveforms is shown in Fig. 5.

VI. BUILDING THE SURROGATE MODEL

Starting from the 104 NR waveforms described in Secs. II and III, we construct hybrid waveforms as described in Sec. V. In this section we describe our method to construct a surrogate model for these hybrid waveforms.

A. Processing the training data

Before building a surrogate model, we process the hybrid waveforms as follows.

1. Time shift

We shift the time arrays of the hybrid waveforms such that the peak of the total amplitude

$$A_{\rm tot} = \sqrt{\sum_{l,m} |\mathfrak{h}_{\ell m}|^2} \tag{38}$$

occurs at $t = 0$ for each waveform.

2. Frequency and mass ranges of validity

The length of a hybrid waveform is set by choosing a starting orbital frequency ω_0 for the inspiral waveform; we use $\omega_0 = 2 \times 10^{-4}$ rad$/M$ for all waveforms. However, for the same starting frequency, the length in time of the waveform is different for different mass ratios and spins. Since we want to construct a time-domain surrogate model, we require a common time array for all hybrid waveforms. The initial time for the surrogate is determined by the shortest hybrid waveform in the training dataset; this waveform begins at a time $\sim 5.4 \times 10^8 M$ before the peak. We truncate all hybrid waveforms to this initial time value.

The largest starting orbital frequency among the truncated hybrid waveforms is $\omega_0 = 2.9 \times 10^{-4}$ rad$/M$, which sets the low-frequency limit of validity of the surrogate. For LIGO, assuming a starting GW frequency of 20 Hz, the (2, 2) mode of the surrogate is valid for total masses $M \geq 0.9\ M_\odot$. The highest spin-weighted spherical harmonic mode we include in the surrogate model is (5, 5), for which the frequency is $5/2$ times that of the (2, 2) mode. Therefore, all modes of the surrogate are valid for $M \geq 2.25\ M_\odot$. This coverage of total mass is sufficient to model all stellar mass binaries of interest for ground-based detectors; for an equal-mass binary neutron star system, the total mass is $M \sim 2.7\ M_\odot$.

3. Downsampling and common time samples

Because the hybrid waveforms are so long, it is not practical to sample the entire waveform with the same step size we use for the NR waveforms ($0.1M$). Fortunately, the early low-frequency portion of each waveform requires

sparser sampling than the later high-frequency portion. We therefore downsample the time arrays of the truncated hybrid waveforms to a common set of time samples. We choose these samples so that there are five points per orbit for the above-mentioned shortest hybrid waveform in the training dataset, except for $t \geq -1000M$ we choose uniform time samples separated by $0.1M$. This ensures that we have a denser sampling rate at late times when the frequency is higher. We retain times up to $135M$, which is sufficient to capture the entire ringdown.

Before downsampling, we first transform the waveform into the co-orbital frame, defined as

$$\mathfrak{h}^{C}_{\ell m} = \mathfrak{h}_{\ell m} e^{\mathrm{i}m\phi_{\rm orb}}, \tag{39}$$

$$\mathfrak{h}_{22} = A_{22} e^{-\mathrm{i}\phi_{22}}, \tag{40}$$

$$\phi_{\rm orb} = \frac{\phi_{22}}{2}, \tag{41}$$

where $\mathfrak{h}_{\ell m}$ is the inertial frame hybrid waveform, $\phi_{\rm orb}$ is the orbital phase, and ϕ_{22} is the phase of the (2, 2) mode. The co-orbital frame can be thought of as roughly corotating with the binary, since we perform a time-dependent rotation given by the instantaneous orbital phase. Therefore the waveform is a slowly varying function of time in this frame, increasing the accuracy of interpolation to the chosen common time samples. For the (2, 2) mode we save the downsampled amplitude A_{22} and phase ϕ_{22}, while for all other modes we save $\mathfrak{h}^{C}_{\ell m}$. We find that this downsampling results in interpolation errors $\mathcal{E} \lesssim 10^{-10}$ [defined in Eq. (18)] for all hybrid waveforms.

4. Phase alignment

After downsampling to the common temporal grid of the surrogate, we rotate the waveforms about the z axis such that the orbital phase $\phi_{\rm orb}$ is zero at $t = -1000M$. Note that this by itself would fix the physical rotation up to a shift of π. When generating the inspiral waveforms for hybridization, we align the system such that the heavier BH is on the positive x axis at the initial frequency; this fixes the π ambiguity. Therefore, after this phase rotation, the heavier BH is on the positive x axis at $t = -1000M$ for all waveforms.[6]

B. Decomposing the data

It is much easier to build a model for slowly varying functions of time. Therefore, rather than work with the inertial frame strain $\mathfrak{h}_{\ell m}$, which is oscillatory, we work with simpler "waveform data pieces," as explained below. We build a separate surrogate for each waveform data piece. When evaluating the full surrogate model, we first evaluate the surrogate of each data piece and then recombine the data pieces to get the inertial frame strain.

A common choice in literature when working with nonprecessing waveforms has been to decompose the complex strain into an amplitude and phase, each of which is a slowly varying function of time,

$$\mathfrak{h}_{\ell m} = A_{\ell m} e^{-\mathrm{i}\phi_{\ell m}}. \tag{42}$$

However, when $q = 1$ and $\chi_{1z} = \chi_{2z}$, the amplitude of odd-m modes becomes zero due to symmetry. This means that the phase becomes meaningless, so one has to treat such cases separately. For example, Ref. [26] used specialized basis functions for the odd-m modes that captured the divergent behavior of the phase in the equal-mass limit.

To avoid this issue, instead of using the amplitude and phase we use the real and imaginary parts of the co-orbital frame strain $\mathfrak{h}^{C}_{\ell m}$, defined in Eq. (39), for all nonquadrupole modes. The co-orbital frame strain is always meaningful: in the special, symmetric case mentioned above, the co-orbital frame strain for the odd-m modes just goes to zero, rather than diverge. For the (2, 2) mode we use the amplitude[7] A_{22} and phase ϕ_{22}.

As mentioned above, our hybrid waveforms are very long, typically containing $\sim 3 \times 10^4$ orbits. This presents new challenges that are not present for pure-NR surrogates. For instance, ϕ_{22} sweeps over $\sim 4 \times 10^5$ rad for a typical hybrid waveform. We find that the accuracy of the surrogate model at early times improves if we first subtract a PN-derived approximation to the phase, model the phase difference rather than ϕ_{22}, and then add back the PN contribution when evaluating the surrogate model. In particular, we use the leading-order TaylorT3 approximant [97]. For this approximant, the phase is given as an analytic, closed-form function of time. Therefore, even though TaylorT3 is known to be less accurate than some other approximants [98], its speed makes it ideal for our purpose as we only need it to capture the general trend. At leading order, the TaylorT3 phase is given by

$$\phi^{T3}_{22} = \phi^{T3}_{\rm ref} - \frac{2}{\eta\theta^5}, \tag{43}$$

where $\phi^{T3}_{\rm ref}$ is an arbitrary integration constant, $\theta = [\eta(t_{\rm ref} - t)/(5M)]^{-1/8}$, $t_{\rm ref}$ is an arbitrary time offset, and η is the symmetric mass ratio. Note that ϕ^{T3}_{22} diverges at $t = t_{\rm ref}$. We choose $t_{\rm ref} = 1000M$, long after the end of the waveform (recall that the peak is at $t = 0$), to ensure that we are always far away from this divergence. We choose $\phi^{T3}_{\rm ref}$

[6]Here the BH positions at $t = -1000M$ are defined from the waveform at future-null infinity, using a phase rotation relative to the early inspiral where the BH positions are well defined in PN theory; these positions do not necessarily correspond to the (gauge-dependent) coordinate BH positions in the NR simulation.

[7]Note that for the (2, 2) mode $A_{22} = \mathfrak{h}^{C}_{22}$.

such that $\phi_{22}^{T3} = 0$ at $t = -1000M$; this is the same time at which we align the hybrid phase in Sec. VI A 4.

Instead of modeling ϕ_{22}, we model the residual

$$\phi_{22}^{\rm res} = \phi_{22} - \phi_{22}^{T3}, \tag{44}$$

after removing the leading-order contribution ϕ_{22}^{T3}. By construction, $\phi_{22}^{\rm res}$ goes to zero at $t = -1000M$. We find that, after removing the leading-order TaylorT3 phase, the scale of $\phi_{22}^{\rm res}$ for a typical hybrid is $\sim 10^3$ rad, compared to $\sim 4 \times 10^5$ rad for ϕ_{22}. In essence, this captures almost all of the phase evolution in the early inspiral, simplifying the problem of modeling the phase to the same as modeling the phase of late-inspiral NR waveforms. We stress that the exact form of ϕ_{22}^{T3} (or its physical meaning) is not important, as long as it captures the general trend, since we add the exact same ϕ_{22}^{T3} to our model of $\phi_{22}^{\rm res}$ when evaluating the surrogate. In fact, we find that adding higher-order PN terms in Eq. (43) does not improve the accuracy of the surrogate.

To summarize, we decompose the hybrid waveforms into the following waveform data pieces, each of which is a smooth, slowly varying function of time: $(A_{22}, \phi_{22}^{\rm res})$ for the (2, 2) mode, and the real and imaginary parts of $\mathfrak{h}_{\ell m}^{C}$ for all other modes.[8]

C. Building the surrogate

Once we have the waveform data pieces, we build a surrogate model for each data piece using the procedure outlined in Refs. [22,27], which we only briefly describe here. Note that the steps below are applied independently for each waveform data piece.

1. Greedy basis

We first construct a greedy reduced basis [99] such that the projection errors [cf. Eq. (5) of Ref. [27]] for the entire dataset onto this basis are below a given tolerance. For the basis tolerances we use 10^{-2} rad for the $\phi_{22}^{\rm res}$ data piece, 2×10^{-5} for A_{22}, and 8×10^{-6} for all other data pieces. These are chosen through visual inspection of the basis functions to ensure they are not noisy and based on the expected truncation error of the NR waveforms. For instance, we expect the error in phase to be about 10^{-2} rad.

The greedy procedure is initialized with a single basis function as described in Ref. [27]. Then, at each step in the greedy procedure, the waveform with the highest projection error onto the current basis is added to the basis. Previous work has shown that the resulting greedy reduced basis is robust to different choices of initialization [100]. When computing the basis projection errors, we only include data up to $50M$ after the peak. We find that this helps avoid noisy basis functions. This is particularly important for the phase data piece as this becomes meaningless at late times, when the waveform amplitude becomes very small.

2. Empirical interpolation

Next, using a different greedy procedure, we construct an empirical interpolant [101–103] in time. This picks out the most representative time nodes, where the number of time nodes is the same as the number of greedy basis functions. We require that the start of the waveform always be included as a time node for all data pieces. This is a useful modeling choice because the magnitude of the waveform data pieces in the very early inspiral can be smaller than the basis tolerances mentioned above. By requiring the first index to be an empirical time node, we enforce an anchor point that ensures the waveform data piece has the right magnitude at the start of the waveform. Furthermore, we do not allow any empirical time nodes at times $> 50M$, since we expect this part to be dominated by noise (especially for the phase data piece).

3. Parametric fits

Finally, for each time node, we construct a fit across the parameter space. The fits are done using the Gaussian process regression (GPR) fitting method described in the supplemental material of Ref. [104]. Following Ref. [104], we parametrize our fits using $\log(q)$, $\hat{\chi}$, and χ_a. Here $\hat{\chi}$ is the spin parameter entering the GW phase at leading order [16,105–107] in the PN expansion,

$$\chi_{\rm eff} = \frac{q\chi_{1z} + \chi_{2z}}{1+q}, \tag{45}$$

$$\hat{\chi} = \frac{\chi_{\rm eff} - 38\eta(\chi_{1z} + \chi_{2z})/113}{1 - 76\eta/113}, \tag{46}$$

and χ_a is the "antisymmetric spin,"

$$\chi_a = \frac{1}{2}(\chi_{1z} - \chi_{2z}). \tag{47}$$

The fit accuracy, and as a result the accuracy of the surrogate model, improves noticeably when using $\log(q)$, compared to q or η.

D. Evaluating the surrogate

When evaluating the surrogate waveform, we first evaluate each surrogate waveform data piece. Next, we compute the phase of the (2, 2) mode,

$$\phi_{22}^{\rm S} \equiv \phi_{22}^{\rm res,S} + \phi_{22}^{T3}, \tag{48}$$

where $\phi_{22}^{\rm res,S} \approx \phi_{22}^{\rm res}$ is the surrogate model for $\phi_{22}^{\rm res}$ and ϕ_{22}^{T3} is given in Eq. (43). If the waveform is required at a uniform

[8]For $m = 0$ modes of nonprecessing systems, $\mathfrak{h}_{\ell m}^{C}$ is purely real (imaginary) for even (odd) ℓ, so we ignore the imaginary (real) part for these modes.

sampling rate, we interpolate each waveform data piece from the sparse time samples used to construct the model to the required time samples, using a cubic-spline interpolation scheme. Finally, we use Eqs. (39)–(41) to reconstruct the surrogate prediction for the inertial frame strain.

VII. RESULTS

In order to estimate the difference between two waveforms, $\mathfrak{h}_1$ and $\mathfrak{h}_2$, we use the mismatch, defined in Eq. (2), but in this section instead of Eq. (3) we use the frequency-domain inner product

$$\langle \mathfrak{h}_1, \mathfrak{h}_2 \rangle = 4\mathrm{Re} \int_{f_{\min}}^{f_{\max}} \frac{\tilde{\mathfrak{h}}_1(f)\tilde{\mathfrak{h}}_2^*(f)}{S_n(f)} df, \tag{49}$$

where $\tilde{\mathfrak{h}}(f)$ indicates the Fourier transform of the complex strain $\mathfrak{h}(t)$, $*$ indicates a complex conjugation, Re indicates the real part, and $S_n(f)$ is the one-sided power spectral density of a GW detector. We taper the time-domain waveform using a Planck window [108] and then zero pad to the nearest power of 2. We further zero pad the waveform to increase the length by a factor of 8 before performing the Fourier transform. The tapering at the start of the waveform is done over 1.5 cycles of the (2, 2) mode. The tapering at the end is done over the last $20M$. Note that our model contains times up to $135M$ after the peak of the waveform amplitude, and the signal has essentially died down by the last $20M$.

We compute mismatches following the procedure described in Appendix D of Ref. [27]: the mismatches are optimized over shifts in time, polarization angle, and initial orbital phase. Both plus and cross polarizations are treated on an equal footing by using a two-detector setup where one detector sees only the plus and the other only the cross polarization. We compute the mismatches at 37 points uniformly distributed on the sky in the source frame, and we use all available modes of a given waveform model.

When computing flat noise mismatches ($S_n = 1$), we take $f_{\min}$ to be the frequency of the (2, 2) mode at the end of the initial tapering window, and $f_{\max} = 5f_{22}^{\mathrm{peak}}$, where f_{22}^{peak} is the frequency of the (2, 2) mode at its peak. This choice of $f_{\max}$ ensures that we capture the peak frequencies of all modes considered in this work, including the (5, 5) mode, whose frequency has the highest multiple of the (2, 2) mode frequency of all the modes we model. We also compute mismatches with the advanced LIGO design sensitivity zero-detuned-high-power noise curve [109] with $f_{\min} = 20$ and $f_{\max} = 2000$ Hz.

A. Surrogate errors

We evaluate the accuracy of our new surrogate model, NRHybSur3dq8, by computing mismatches against hybrid waveforms. For this, we compute "out-of-sample" errors as follows. We first randomly divide the 104 training waveforms into groups of ∼5 waveforms each. For each group, we build a trial surrogate using the remaining ∼99 training waveforms and test against these five validation ones. We also compute the mismatch between an existing higher-mode waveform model, SEOBNRv4HM [15], and the hybrid waveforms.

Figure 6 summarizes mismatches of both NRHybSur3dq8 and SEOBNRv4HM versus the hybrid waveforms. We use all available modes for each waveform model. In the left panel we show mismatches computed using a flat noise

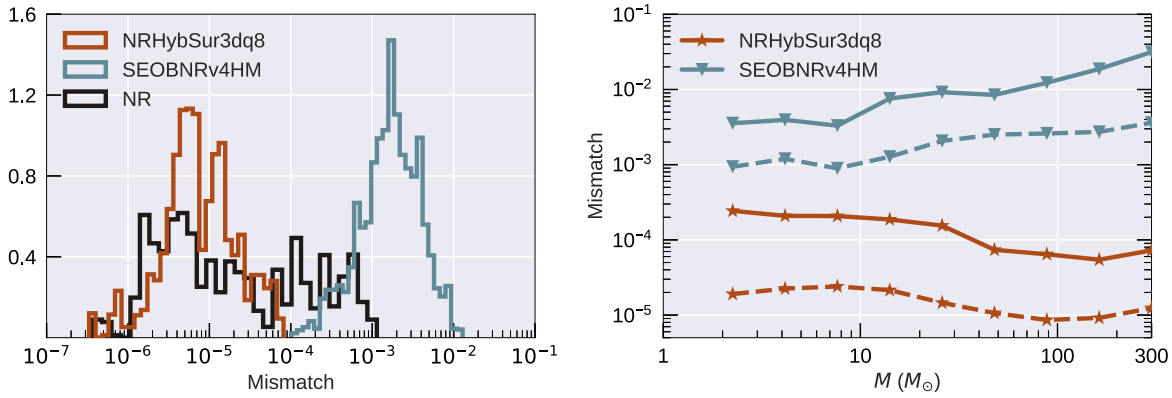

FIG. 6. Errors in NRHybSur3dq8 and SEOBNRv4HM when compared against hybrid waveforms. For NRHybSur3dq8, we show out-of-sample errors. Mismatches are computed at several points in the sky of the source frame using all available modes in each waveform: For the hybrid waveforms and NRHybSur3dq8, that is $\ell \leq 4$ and (5, 5), but not (4, 1) or (4, 0). For SEOBNRv4HM that is (2, 2),(2, 1), (3, 3),(4, 4), and (5, 5). (Left) Mismatches computed using a flat noise curve, but including only the late inspiral part of the waveforms, starting at $-3500M$ before the peak. Therefore, we are essentially comparing only to the NR part of the hybrid waveforms. For comparison, we also show the NR resolution error, obtained by comparing the two highest available resolutions. The histograms are normalized such that the area under each curve is 1 when integrated over $\log_{10}(\mathrm{Mismatch})$. (Right) Mismatches as a function of total mass, computed using the advanced LIGO design sensitivity noise curve. Here we compare against the full hybrid waveforms. The solid (dashed) lines show the 95th percentile (median) mismatch values over points on the sky as well as different hybrid waveforms.

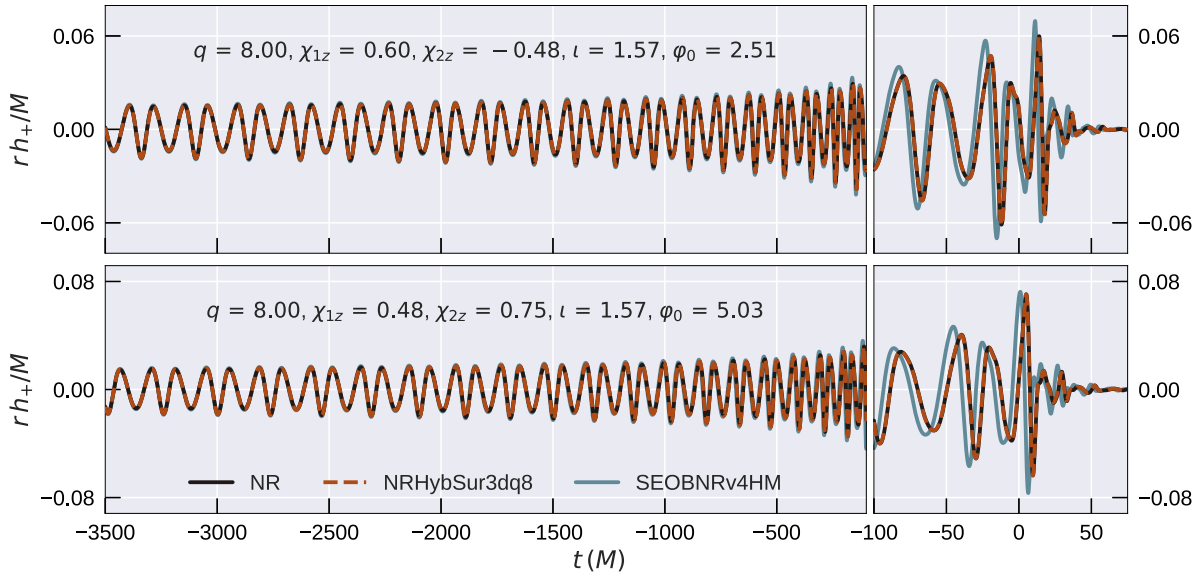

FIG. 7. The plus polarization of the waveforms for the cases that result in the largest mismatch for NRHybSur3dq8 (top) and SEOBNRv4HM (bottom) in the left panel of Fig. 6. We also show the corresponding hybrid waveforms (labeled as NR because only the late part is shown). Each waveform is projected using all available modes for that model, along the direction which results in the largest mismatch for NRHybSur3dq8 (SEOBNRv4HM) in the top (bottom) panel. Note that NRHybSur3dq8 is evaluated using trial surrogates that are not trained using these cases. The binary parameters and the direction in the source frame are indicated in the inset text. All waveforms are time shifted such that the peak of the total waveform amplitude occurs at $t = 0$ [using all available modes, according to Eq. (38)]. Then the waveform modes are rotated about the z axis such that the orbital phase is zero at $t = -3500M$.

curve over the NR part of the hybrid waveforms (to do this, we truncate the waveforms and begin tapering at $t = -3500M$). We see that the mismatches for NRHybSur3dq8 are about 2 orders of magnitude lower than that of SEOBNRv4HM. We compare this with the truncation error in the NR waveforms themselves, by computing the mismatch between the two highest available resolutions of each NR waveform. The errors in the surrogate model are well within the truncation error of the NR simulations. Note that NR error estimated in this manner is a conservative estimate; if we treat the high-resolution simulation as the fiducial case, the NR curve in Fig. 6 can be thought of as the error in the lower-resolution simulation. This explains why the errors in the surrogate are smaller than the NR errors. We suspect that the error of the high-resolution simulations is close to the surrogate model's error.

The right panel of Fig. 6 shows mismatches computed using the advanced LIGO design sensitivity noise curve. The mismatches are now dependent on the total mass of the system, so we show mismatches for masses starting at the lower limit of the range of validity of the surrogate: $M \geq 2.25\ M_\odot$. The 95th percentile mismatches for NRHybSur3dq8 are always below $\sim 3 \times 10^{-4}$ in the mass range $2.25\,M_\odot \leq M \leq 300\,M_\odot$. At high masses ($M \gtrsim 40\,M_\odot$), where the merger and ringdown are more prominent, our model is more accurate than SEOBNRv4HM by roughly 2 orders of magnitude, in agreement with the left panel of Fig. 6.

For high masses only the last few orbits of the hybrid waveforms are in the LIGO band, and the hybrid waveforms are effectively the same as the NR waveforms. For low masses, the errors in the right panel of Fig. 6 quantify how well different models reproduce the hybrid waveforms. However, this comparison cannot account for the errors in the hybridization procedure itself. We provide some evidence for the fidelity of the hybrid waveforms in Sec. VII B, by comparing against some long NR waveforms.

Figure 7 shows NRHybSur3dq8 and SEOBNRv4HM waveforms for the cases leading to the largest errors in the left panel of Fig. 6. The surrogate shows very good agreement with the NR waveform, even for its worst case. SEOBNRv4HM shows a noticeably larger deviation that cannot all be accounted for with a time and/or phase shift. Note that we align the time and orbital phase of the waveforms in Fig. 7.

We note that the main improvement over SEOBNRv4HM is not due to the inclusion of more modes. We find that the agreement between SEOBNRv4HM and the NR (Hybrid) waveforms in Figs. 6 and 7 improves only marginally when restricting the NR-hybrid waveforms to the same set of modes as SEOBNRv4HM.

B. Hybridization errors

The errors described in Sec. VII A are computed by comparing the surrogate against hybrid waveforms, hence

they do not include the errors in the hybridization procedure or the errors from EOB-corrected PN waveforms (cf. Sec. IV B) we use for the early inspiral. To estimate these errors, we compare the surrogate against a few very long NR simulations.[9] We perform five new simulations that are $\sim 10^5 M$ long and two that are $\sim 3 \times 10^4 M$ long. These have been assigned the identifiers SXS:BBH:1412–SXS:BBH:1418 and will be made publicly available in the upcoming update of the SXS public catalog [58]. In addition, we use two simulations of length $\sim 3 \times 10^4 M$ from Ref. [111]. These nine simulations are represented as square markers in Fig. 2 and have not been used in training the surrogate. The surrogate was trained against hybrid waveforms whose NR duration varied between $3270M$ and $4227M$. Therefore, comparing against long NR waveforms, which include the early inspiral, is a good way to estimate the hybridization error.

We begin by repeating the mismatch computation from the right panel of Fig. 6, using the $10^5 M$ long NR waveforms. This is shown in Fig. 8. We also show the errors in the NR simulations, estimated by comparing the two highest available NR resolutions. We find that the mismatches between the surrogate and the long NR waveforms for $M > 30\, M_\odot$ are below 10^{-4}, in agreement with Fig. 6. For lower masses, the mismatches quickly increase and can be as high as $\sim 10^{-2}$. However, this increase in mismatch is accompanied by an increase in the error of the NR waveforms. This is expected, since for very long NR waveforms the accumulated phase error is a dominant source of numerical error, which becomes increasingly relevant for low mass systems as more of the waveform moves in-band. Therefore, in Fig. 8, at low masses, the comparison between the surrogate and NR waveforms is largely dominated by the numerical resolution error of the long NR waveforms themselves.

We find that a better test of the hybridization procedure, one that is less sensitive to NR phase accumulation errors, is to compare against different segments of the NR waveform. Since the phase errors accumulate over a large number of cycles, by looking at smaller segments we ensure that this contribution is not the dominant error. To be precise, we compare the surrogate and the NR data, using segments of length $\Delta t = 5 \times 10^3 M$ ending at a particular number of orbits before the peak of the waveform. For each segment we compute mismatches at several points in the sky using a flat noise curve. By varying the number of orbits to the peak, we can cover the entire NR waveform including the early inspiral region where the surrogate depends on the hybridization procedure. These errors are shown in Fig. 9. We find that, in each segment, the mismatch between the surrogate and the NR data is, in general, lower or comparable to the NR resolution error. Therefore, the surrogate reproduces the NR data accurately in the early inspiral and the hybridization errors are smaller than or comparable to the NR resolution error for these cases. We note that the surrogate errors in Fig. 9 depend on the length of the segment considered and are only meaningful when compared to the NR errors in the same segment.

Unfortunately, long NR simulations such as these are not available at regions of the parameter space where both mass ratio and spin magnitudes are large. These are the cases where PN is expected to perform poorly, so we expect larger hybridization errors for these cases.

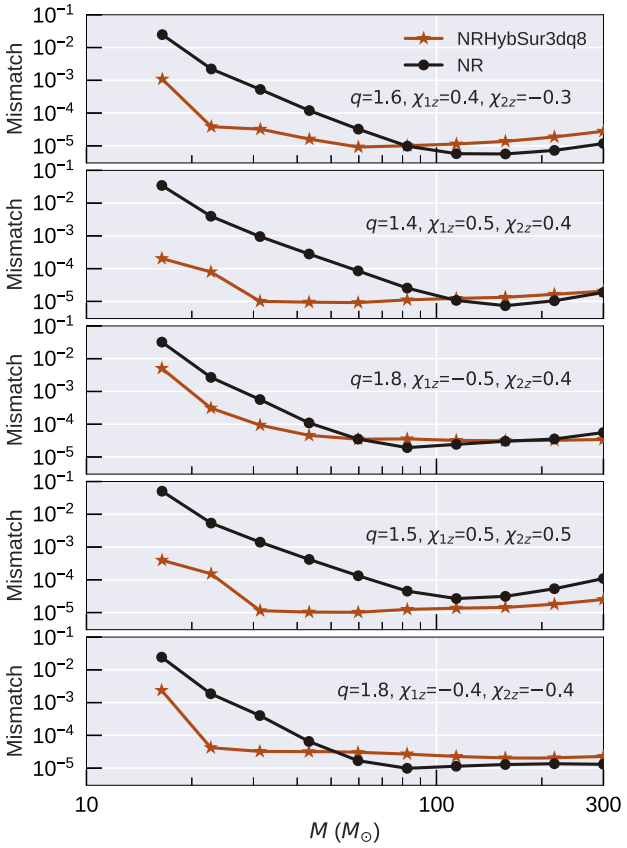

FIG. 8. Comparisons between the NRHybSur3dq8 surrogate model and a few NR waveforms of $\sim 10^5 M$ in duration. We also show the NR resolution error. The 95th percentile mismatches (over points in the sky) are shown as a function of total mass. (Inset) Indicates the mass ratio and component spins. Mismatches are computed using the advanced LIGO design sensitivity noise curve. To best assess the error introduced by the hybridization procedure we use the same set of modes for the NR waveforms as the surrogate. At low masses, the hybridization errors (red circles) become less reliable measures of accuracy due to the large NR resolution error (black circles) itself. Figure 9 describes a refined comparison to improve the assessment at low masses.

[9]Note that, for these long NR simulations, the outer boundary location is chosen based on the length of the simulations [59] so as to avoid unphysical center-of-mass accelerations seen in earlier long-duration runs [110].

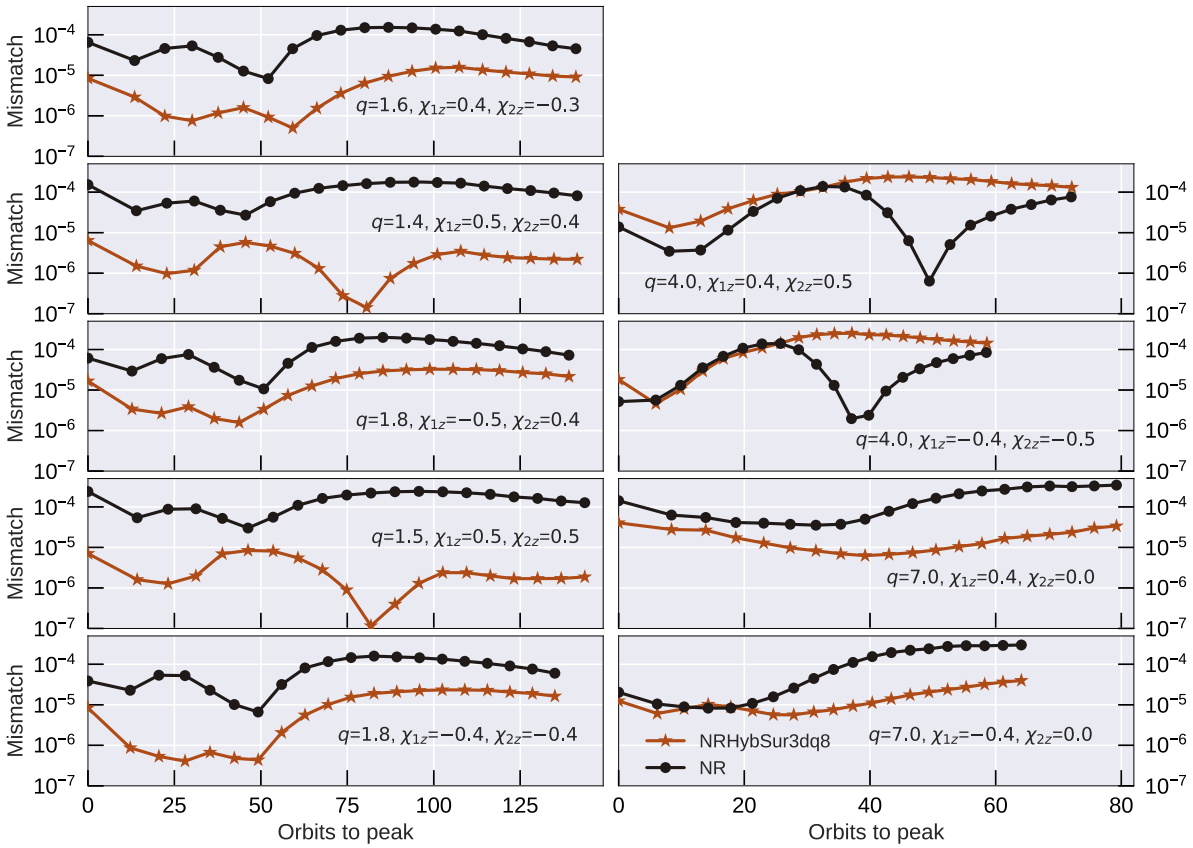

FIG. 9. Errors in the NRHybSur3dq8 surrogate model against long NR waveforms, but only looking at segments of length $\Delta t = 5 \times 10^3 M$ individually. Each point represents one segment that ends at a specified number of orbits before the waveform peak, as plotted on the horizontal axis. Therefore, going from left to right in the figure, we plot segments that start earlier in the inspiral. We also show the NR resolution error in the same segments. (Inset) Indicates the mass ratio and component spins. We show 95th percentile mismatches (over points in the sky), computed using a flat noise curve. We use the same set of modes for the NR waveforms as the surrogate. We find that, in general, the surrogate error is lower than or comparable to the NR resolution error throughout the inspiral.

C. Extrapolation outside the training range

We now investigate the efficacy of NRHybSur3dq8 to extrapolate beyond its training parameter range by comparing against SpEC NR simulations [58,111–114] at larger mass ratios ($8 < q \leq 10$) and/or larger spin magnitudes ($|\chi_{1z}| > 0.8$ or $|\chi_{12}| > 0.8$). These NR simulations are represented as triangle markers in Fig. 2.

Figure 10 shows mismatches for NRHybSur3dq8 when compared against these simulations. We find that the surrogate extrapolates remarkably well, with the mismatch always $\lesssim 4 \times 10^{-4}$ for all cases, which include mass ratios up to $q = 10$ and spin magnitudes up to $|\chi| = 0.998$. However, the extrapolation errors can be about half an order of magnitude larger than errors within the training range. Note that NR simulations with both high mass ratios and high spin magnitudes are not currently available, and the ones used here represent the most extreme cases found in the SXS catalog. We do not hybridize these simulations before comparing to NRHybSur3dq8 because several of them are too short. In Fig. 10, the minimum mass for each case is chosen to be the lowest mass at which all used modes of the NR simulation lie fully in the LIGO band with a low-frequency cutoff of 20 Hz.

At much higher mass ratios than those tested here, such as $q = 15$, we find that the waveforms generated by the surrogate can have "glitches" in the time series. Therefore, we recommend the surrogate be used for $q \leq 10$ and $|\chi_{1z}|, |\chi_{2z}| \leq 1$. However, we advise caution with any extrapolation in general.

D. Mode mixing

Numerical relativity waveforms are extracted as spin-weighted spherical harmonic modes [33,34]. However, in the ringdown regime, the natural basis to use is the spin-weighted *spheroidal* harmonic basis [115,116]. A spherical harmonic mode $\mathfrak{h}_{\ell m}$ can be written as a linear combination of all spheroidal harmonic modes $\mathfrak{h}^S_{\ell m}$ with the same m index [117]. Therefore, during the ringdown, we expect leakage of power between different spherical harmonic modes with the same m. This is referred to as mode mixing.

Since the surrogate accurately reproduces the spherical harmonic modes from the NR simulations, it also captures

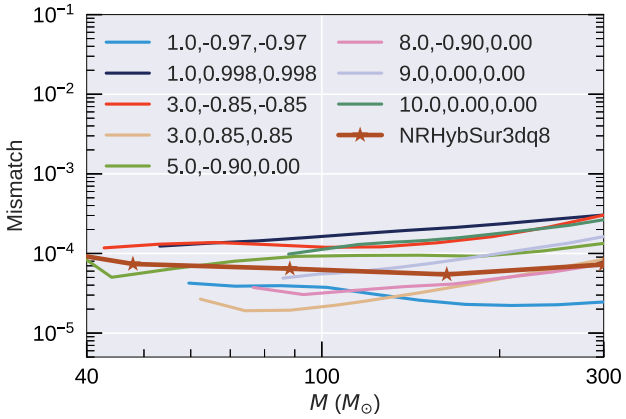

FIG. 10. Errors in NRHybSur3dq8 when evaluated outside its training range. The 95th percentile mismatches (over points in the sky) are shown as a function of total mass for different extrapolated cases. These are computed using the advanced LIGO design sensitivity noise curve. To best assess the error introduced by the extrapolation, we use the same set of modes for the NR waveforms as the surrogate. The labels indicate the mass ratio and component spins (q, χ_{1z}, χ_{2z}). For comparison we reproduce the 95th percentile mismatches for NRHybSur3dq8 within its training range from the right panel of Fig. 6.

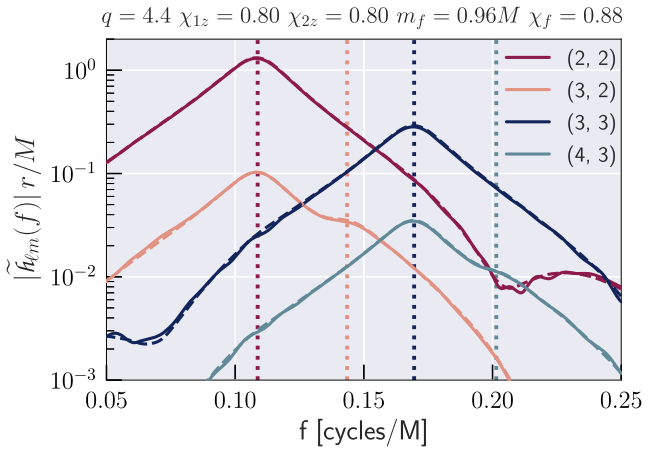

FIG. 11. Mode mixing between spherical harmonic modes is clearly seen in the ringdown signal of the NR waveform and is accurately reproduced by the surrogate. The absolute values of the Fourier transform of different spherical harmonic modes are shown as solid (dashed) curves for the surrogate (NR). The dotted vertical lines indicate the frequencies of the fundamental QNM overtone of these modes. The component parameters as well as the remnant mass and spin are shown in the text above the figure.

this mode mixing. We demonstrate this for an example case in Fig. 11. Here we compute the Fourier transform of different spherical harmonic modes in the ringdown stage of the waveform. Before computing the Fourier transform, we first drop all data before $t = 20M$, where $t = 0$ corresponds to the peak of the waveform amplitude [cf. Eq. (38)]. Then, we taper the data between $t = 20M$ and $t = 40M$, as well as the last $10M$ of the time series, using a Planck window [108]. The tapering width at the start is chosen such that the remaining signal is dominated by the fundamental quasinormal mode (QNM) overtone. Figure 11 shows the absolute value of these Fourier transforms for different modes, for both the surrogate and the NR waveform. In addition, we show the frequency of the fundamental QNM overtone for each mode [118].

Note that the (2, 2) mode and the (3, 2) mode have the same m index, the condition required for mode mixing. We see that the peak of the (2, 2) mode agrees with the QNM frequency as expected. For the (3, 2) mode, however, while there are features of a peak at the expected QNM frequency, there is a much larger peak at the frequency of the (2, 2) mode. This is because some of the power of the stronger (2, 2) mode has leaked into the (3, 2) mode due to mode mixing. Mode mixing can also be seen for the (3, 3) and (4, 3) modes, which also have the same m index. Figure 11 shows that, not only does the surrogate agree with NR in the ringdown, it also reproduces the mode mixing present in the NR data.

E. Evaluation cost

Figure 12 shows the evaluation cost for NRHybSur3dq8, at different total masses, starting at 20 Hz, and using a sampling rate of 4096 Hz. This suggests that NRHybSur3dq8 is fast enough for direct use in parameter estimation. We also show the evaluation cost per mode. Note that the total cost as well the cost per mode in Fig. 12 include the cost of a fast Fourier transform (FFT). We perform the FFT only once, after summing over all modes in the time domain. This cost is also shown separately in Fig. 12. Finally, we show the evaluation cost of

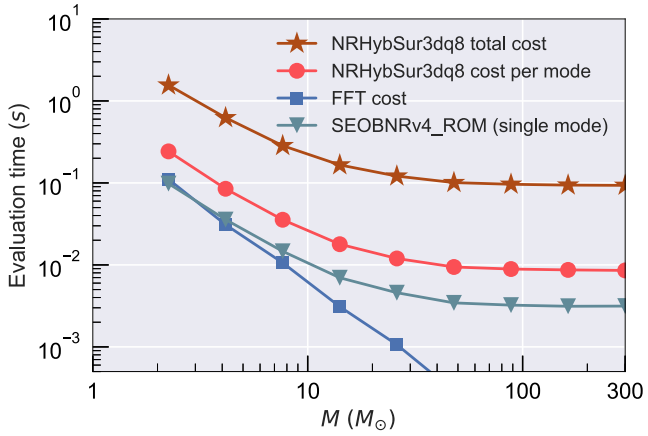

FIG. 12. Evaluation cost for NRHybSur3dq8 including the cost of a FFT. We show the cost for evaluating all 11 modes modeled by NRHybSur3dq8, as well as the cost per mode. The FFT cost is included in both of the above but also shown separately. We also show the evaluation cost of SEOBNRv4_ROM, which includes only the (2, 2) mode. The evaluation cost is computed by averaging over 64 points uniformly distributed in the parameter space, $q \leq 8$ and $|\chi_{1z}|, |\chi_{2z}| \leq 0.8$.

SEOBNRv4_ROM [17], a Fourier domain reduced-order model (ROM) version of SEOBNRv4. Note that SEOBNRv4_ROM models only the (2, 2) mode. Comparing the cost for SEOBNRv4_ROM to the cost per mode of NRHybSur3dq8 suggests that the evaluation cost of NRHybSur3dq8 can be reduced by a factor of ∼2.5 by building a Fourier domain ROM along the lines of Ref. [23].

At low masses, where the waveform is very long, the dominant costs for NRHybSur3dq8 are due to the temporal interpolation from the sparse domain of the surrogate to the required time samples and the FFT. At high masses, where the waveform is short, the interpolation and FFT are cheap and the dominant cost for NRHybSur3dq8 is due to the GPR evaluations for the parametric fits. SEOBNRv4_ROM instead uses tensor spline interpolation for the parametric fits [17], which accounts for the main difference in the evaluation cost per mode at high masses.

These tests were performed on a single core on a 3.1 GHz Intel Core i5 processor. Both NRHybSur3dq8 and SEOBNRv4_ROM were evaluated using a C implementation in the LIGO algorithm library [49]. The PYTHON implementation of NRHybSur3dq8 in *gwsurrogate* [48] is slower than the C implementation by at most a factor of 2.

VIII. CONCLUSION

We present NRHybSur3dq8, the first NR-based surrogate waveform model that spans the entire LIGO bandwidth, valid for stellar mass binaries with total masses $M \geq 2.25\ M_\odot$. This model is trained on 104 NR-PN/EOB hybrid waveforms of nonprecessing quasicircular BBH systems with mass ratios $q \leq 8$ and spin magnitudes $|\chi_{1z}|, |\chi_{2z}| \leq 0.8$. The parametric fits for this model are performed using Gaussian process regression. This model includes the following spin-weighted spherical harmonic modes: $\ell \leq 4$ and (5, 5), but not (4, 1) or (4, 0). We make our model available publicly through the easy-to-use PYTHON package *gwsurrogate* [48]. In addition, our model is implemented in C with PYTHON wrapping in the LIGO algorithm library [49]. We provide an example PYTHON evaluation code in [50].

Through a cross-validation study, we show that the surrogate accurately reproduces the hybrid waveforms. The mismatch between them is always less than $\sim 3 \times 10^{-4}$ for total masses $2.25\ M_\odot \leq M \leq 300\ M_\odot$. For high masses ($M \gtrsim 40\ M_\odot$), where the merger and ringdown are more prominent, we show roughly a 2 orders of magnitude improvement over the current state-of-the-art model with nonquadrupole modes, SEOBNRv4HM [15].

By comparing against several long NR simulations, we show that the errors in our hybridization procedure are comparable or lower than the resolution error in current NR simulations. In addition, by comparing against available NR simulations at higher mass ratios and spins, we show that our model extrapolates reasonably well outside its training range. Based on these tests, we are cautiously optimistic that the surrogate can be used for $q \leq 10$ and $|\chi_{1z}|, |\chi_{2z}| \leq 1$, and we leave a more detailed investigation for future work.

A. Future work

While our tests of the hybridization procedure are encouraging, long NR simulations are available only for low mass ratios and low spin magnitudes. Therefore, we have no means to test hybridization at high mass ratios and/or high spins, where the PN model is expected to perform poorly. An improved surrogate model and refined study of the hybridization errors will require longer inspiral waveforms with greater coverage of the parameter space.

Another extension of interest is towards larger mass ratios and spin magnitudes. While the surrogate extrapolates very well when compared to available simulations at larger mass ratios and spins, no NR simulations are available with both large mass ratios ($q > 8$) and large spins ($\chi > 0.8$). Therefore, our model is untested in that region of parameter space and it might be necessary to add training points there. The model could also be extended to include precession and/or eccentricity, however, this is more challenging because of the enlarged parameter space as well as more complicated hybridization.

Finally, as mentioned in Sec. VII E, the evaluation time of NRHybSur3dq8 can likely be reduced by constructing a Fourier domain ROM [23] of the time-domain model.

We leave these explorations to future work.

ACKNOWLEDGMENTS

We thank Matt Giesler for helping carry out the new SPEC simulations used in this work. We thank Michael Boyle, Kevin Barkett, Matt Giesler, Yanbei Chen, and Saul Teukolsky for useful discussions. We thank Patricia Schmidt for careful and detailed feedback on an earlier draft of this manuscript. V. V. and M. S. are supported by the Sherman Fairchild Foundation and NSF Grants No. PHY–170212 and No. PHY–1708213 at Caltech. L. E. K. acknowledges support from the Sherman Fairchild Foundation and NSF Grant No. PHY-1606654 at Cornell. S. E. F. is partially supported by NSF Grant No. PHY-1806665. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation Grant No. ACI-1548562. This research is part of the Blue Waters sustained-petascale computing project, which is supported by the National Science Foundation (Grants No. OCI-0725070 and No. ACI-1238993) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications. Simulations were performed on NSF/NCSA Blue Waters under allocation

NSF PRAC–1713694; on the Wheeler cluster at Caltech, which is supported by the Sherman Fairchild Foundation and by Caltech; and on XSEDE resources Bridges at the Pittsburgh Supercomputing Center, Comet at the San Diego Supercomputer Center, and Stampede and Stampede2 at the Texas Advanced Computing Center, through allocation TG-PHY990007N. Computations for building the model were performed on Wheeler and Stampede2.

[1] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Observation of Gravitational Waves from a Binary Black Hole Merger, Phys. Rev. Lett. **116**, 061102 (2016).
[2] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral, Phys. Rev. Lett. **119**, 161101 (2017).
[3] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence, Phys. Rev. Lett. **116**, 241103 (2016).
[4] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2, Phys. Rev. Lett. **118**, 221101 (2017).
[5] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), GW170608: Observation of a 19-solar-mass binary black hole coalescence, Astrophys. J. **851**, L35 (2017).
[6] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), GW170814: A Three-Detector Observation of Gravitational Waves from a Binary Black Hole Coalescence, Phys. Rev. Lett. **119**, 141101 (2017).
[7] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), GWTC-1: A gravitational-wave transient catalog of compact binary mergers observed by LIGO and Virgo during the first and second observing runs, arXiv:1811.12907.
[8] J. Aasi *et al.* (LIGO Scientific Collaboration), Advanced LIGO, Classical Quantum Gravity **32**, 074001 (2015).
[9] F. Acernese *et al.* (Virgo Collaboration), Advanced Virgo: A second-generation interferometric gravitational wave detector, Classical Quantum Gravity **32**, 024001 (2015).
[10] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Prospects for observing and localizing gravitational-wave transients with advanced Ligo and advanced Virgo, Living Rev. Relativity **19**, 1 (2016).
[11] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), Binary black hole population properties inferred from the first and second observing runs of advanced LIGO and advanced Virgo, arXiv:1811.12940.
[12] 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, arXiv:1809.10113.
[13] Yi Pan, A. Buonanno, A. Taracchini, L. E. Kidder, A. H. Mrou, H. P. Pfeiffer, M. A. Scheel, and B. Szilgyi, Inspiral-merger-ringdown waveforms of spinning, precessing black-hole binaries in the effective-one-body formalism, Phys. Rev. D **89**, 084006 (2014).
[14] L. London, S. Khan, E. Fauchon-Jones, C. Garca, M. Hannam, S. Husa, X. Jimnez-Forteza, C. Kalaghatgi, F. Ohme, and F. Pannarale, First Higher-Multipole Model of Gravitational Waves from Spinning and Coalescing Black-Hole Binaries, Phys. Rev. Lett. **120**, 161102 (2018).
[15] R. Cotesta, A. Buonanno, A. Boh, A. Taracchini, I. Hinder, and S. Ossokine, Enriching the symphony of gravitational waves from binary black holes by tuning higher harmonics, Phys. Rev. D **98**, 084028 (2018).
[16] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Prrer, X. J. Forteza, and A. Boh, Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era, Phys. Rev. D **93**, 044007 (2016).
[17] A. Bohé, L. Shao, A. Taracchini, A. Buonanno, S. Babak, I. W. Harry, I. Hinder, S. Ossokine, M. Pürrer, V. Raymond, T. Chu, H. Fong, P. Kumar, H. P. Pfeiffer, M. Boyle, D. A. Hemberger, L. E. Kidder, G. Lovelace, M. A. Scheel, and B. Szilágyi, 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).
[18] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, A Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms, Phys. Rev. Lett. **113**, 151101 (2014).
[19] A. Taracchini *et al.*, Effective-one-body model for black-hole binaries with generic mass ratios and spins, Phys. Rev. D **89**, 061502 (2014).
[20] Y. Pan, A. Buonanno, M. Boyle, L. T. Buchman, L. E. Kidder, H. P. Pfeiffer, and M. A. Scheel, Inspiral-merger-ringdown multipolar waveforms of nonspinning black-hole binaries using the effective-one-body formalism, Phys. Rev. D **84**, 124052 (2011).
[21] A. K. Mehta, C. K. Mishra, V. Varma, and P. Ajith, Accurate inspiral-merger-ringdown gravitational waveforms for nonspinning black-hole binaries including the effect of subdominant modes, Phys. Rev. D **96**, 124010 (2017).
[22] S. E. Field, C. R. Galley, J. S. Hesthaven, J. Kaye, and M. Tiglio, Fast Prediction and Evaluation of Gravitational Waveforms Using Surrogate Models, Phys. Rev. X **4**, 031006 (2014).
[23] M. Pürrer, Frequency domain reduced order models for gravitational waves from aligned-spin compact binaries, Classical Quantum Gravity **31**, 195010 (2014).
[24] B. D. Lackey, S. Bernuzzi, C. R. Galley, J. Meidam, and C. Van Den Broeck, Effective-one-body waveforms for

binary neutron stars using surrogate models, Phys. Rev. D **95**, 104036 (2017).
[25] Z. Doctor, B. Farr, D. E. Holz, and Michael Prrer, Statistical gravitational waveform models: What to simulate next?, Phys. Rev. D **96**, 123011 (2017).
[26] J. Blackman, S. E. Field, C. R. Galley, B. Szilágyi, M. A. Scheel, M. Tiglio, and D. A. Hemberger, Fast and Accurate Prediction of Numerical Relativity Waveforms from Binary Black Hole Coalescences Using Surrogate Models, Phys. Rev. Lett. **115**, 121102 (2015).
[27] J. Blackman, S. E. Field, M. A. Scheel, C. R. Galley, D. A. Hemberger, P. Schmidt, and R. Smith, A surrogate model of gravitational waveforms from numerical relativity simulations of precessing binary black hole mergers, Phys. Rev. D **95**, 104023 (2017).
[28] J. Blackman, S. E. Field, M. A. Scheel, C. R. Galley, C. D. Ott, M. Boyle, L. E. Kidder, H. P. Pfeiffer, and B. Szilgyi, Numerical relativity waveform surrogate model for generically precessing binary black hole mergers, Phys. Rev. D **96**, 024058 (2017).
[29] E. A. Huerta, C. J. Moore, P. Kumar, D. George, A. J. K. Chua, R. Haas, E. Wessel, D. Johnson, D. Glennon, A. Rebei, A. M. Holgado, J. R. Gair, and H. P. Pfeiffer, Eccentric, nonspinning, inspiral, Gaussian-process merger approximant for the detection and characterization of eccentric binary black hole mergers, Phys. Rev. D **97**, 024031 (2018).
[30] A. J. K. Chua, C. R. Galley, and M. Vallisneri, ROMAN: Reduced-order modeling with artificial neurons, arXiv:1811.05491.
[31] 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).
[32] C. R. Galleyand and P. Schmidt, Fast and efficient evaluation of gravitational waveforms via reduced-order spline interpolation, arXiv:1611.07529.
[33] E. T. Newmanand and R. Penrose, Note on the Bondi–Metzner–Sachs group, J. Math. Phys. (N.Y.) **7**, 863 (1966).
[34] J. N. Goldberg, A. J. Macfarlane, E. T. Newman, F. Rohrlich, and E. C. G. Sudarshan, Spin-*s* spherical harmonics and ð, J. Math. Phys. (N.Y.) **8**, 2155 (1967).
[35] V. Varmaand and P. Ajith, Effects of nonquadrupole modes in the detection and parameter estimation of black hole binaries with nonprecessing spins, Phys. Rev. D **96**, 124024 (2017).
[36] C. Capano, Y. Pan, and A. Buonanno, Impact of higher harmonics in searching for gravitational waves from nonspinning binary black holes, Phys. Rev. D **89**, 102003 (2014).
[37] T. B. Littenberg, J. G. Baker, A. Buonanno, and B. J. Kelly, Systematic biases in parameter estimation of binary black-hole mergers, Phys. Rev. D **87**, 104003 (2013).
[38] J. C. Bustillo, P. Laguna, and D. Shoemaker, Detectability of gravitational waves from binary black holes: Impact of precession and higher modes, Phys. Rev. D **95**, 104038 (2017).
[39] D. A. Brown, P. Kumar, and A. H. Nitz, Template banks to search for low-mass binary black holes in advanced gravitational-wave detectors, Phys. Rev. D **87**, 082004 (2013).
[40] V. Varma, P. Ajith, S. Husa, J. C. Bustillo, M. Hannam, and M. Prrer, Gravitational-wave observations of binary black holes: Effect of nonquadrupole modes, Phys. Rev. D **90**, 124004 (2014).
[41] P. B. Graff, A. Buonanno, and B. S. Sathyaprakash, Missing link: Bayesian detection and measurement of intermediate-mass black-hole binaries, Phys. Rev. D **92**, 022002 (2015).
[42] I. Harry, J. C. Bustillo, and A. Nitz, Searching for the full symphony of black hole binary mergers, Phys. Rev. D **97**, 023004 (2018).
[43] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Effects of waveform model systematics on the interpretation of GW150914, Classical Quantum Gravity **34**, 104002 (2017).
[44] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Directly comparing GW150914 with numerical solutions of Einsteins equations for binary black hole coalescence, Phys. Rev. D **94**, 064035 (2016).
[45] P. Kumar, J. Blackman, S. E. Field, M. Scheel, C. R. Galley, M. Boyle, L. E. Kidder, H. P. Pfeiffer, B. Szilagyi, and S. A. Teukolsky, Constraining the parameters of GW150914 & GW170104 with numerical relativity surrogates, arXiv:1808.08004.
[46] R. O'Shaughnessy, B. Farr, E. Ochsner, H.-S. Cho, V. Raymond, C. Kim, and C.-H. Lee, Parameter estimation of gravitational waves from precessing black hole-neutron star inspirals with higher harmonics, Phys. Rev. D **89**, 102005 (2014).
[47] S. A. Usman, J. C. Mills, and S. Fairhurst, Constraining the inclination of binary mergers from gravitational-wave observations, arXiv:1809.10727.
[48] J. Blackman, S. Field, C. Galley, and Vijay Varma, gwsurrogate, https://pypi.python.org/pypi/gwsurrogate/.
[49] LIGO Scientific Collaboration, LIGO Algorithm Library-LALSuite, free software (GPL) (2018), https://git.ligo.org/lscsoft/lalsuite.
[50] Binary black-hole surrogate waveform catalog, http://www.black-holes.org/surrogates/.
[51] The Spectral Einstein Code, http://www.black-holes.org/SpEC.html.
[52] H. P. Pfeiffer, L. E. Kidder, M. A. Scheel, and S. A. Teukolsky, A multidomain spectral method for solving elliptic equations, Comput. Phys. Commun. **152**, 253 (2003).
[53] G. Lovelace, R. Owen, H. P. Pfeiffer, and T. Chu, Binary-black-hole initial data with nearly-extremal spins, Phys. Rev. D **78**, 084017 (2008).
[54] L. Lindblom, M. A. Scheel, L. E. Kidder, R. Owen, and O. Rinne, A new generalized harmonic evolution system, Classical Quantum Gravity **23**, S447 (2006).
[55] B. Szilágyi, L. Lindblom, and M. A. Scheel, Simulations of binary black hole mergers using spectral methods, Phys. Rev. D **80**, 124010 (2009).
[56] M. A. Scheel, M. Boyle, T. Chu, L. E. Kidder, K. D. Matthews, and H. P. Pfeiffer, High-accuracy waveforms for binary black hole inspiral, merger, and ringdown, Phys. Rev. D **79**, 024003 (2009).
[57] Simulating Extreme Spacetimes, http://www.black-holes.org/.

[58] SXS Collaboration, The SXS collaboration catalog of gravitational waveforms, http://www.black-holes.org/waveforms.

[59] M. Boyle *et al.* (SXS collaboration), The SXS collaboration catalog of gravitational waveforms for merging black holes (to be published).

[60] A. Buonanno, L. E. Kidder, A. H. Mroué, H. P. Pfeiffer, and A. Taracchini, Reducing orbital eccentricity of precessing black-hole binaries, Phys. Rev. D **83**, 104034 (2011).

[61] M. Boyleand and A. H. Mroué, Extrapolating gravitational-wave data from numerical simulations, Phys. Rev. D **80**, 124045 (2009).

[62] M. Boyle, Transformations of asymptotic gravitational-wave data, Phys. Rev. D **93**, 084031 (2016).

[63] M. Boyle, Scri, https://github.com/moble/scri.

[64] L. Santamaría, F. Ohme, P. Ajith, B. Brügmann, N. Dorband, M. Hannam, S. Husa, P. Mösta, D. Pollney, C. Reisswig, E. L. Robinson, J. Seiler, and B. Krishnan, Matching post-Newtonian and numerical relativity waveforms: Systematic errors and a new phenomenological model for non-precessing black hole binaries, Phys. Rev. D **82**, 064016 (2010).

[65] F. Ohme, Analytical meets numerical relativity-status of complete gravitational waveform models for binary black holes, Classical Quantum Gravity **29**, 124002 (2012).

[66] F. Ohme, M. Hannam, and S. Husa, Reliability of complete gravitational waveform models for compact binary coalescences, Phys. Rev. D **84**, 064029 (2011).

[67] I. MacDonald, S. Nissanke, and H. P. Pfeiffer, Suitability of post-Newtonian/numerical-relativity hybrid waveforms for gravitational wave detectors, Classical Quantum Gravity **28**, 134002 (2011).

[68] I. MacDonald, A. H. Mroué, H. P. Pfeiffer, M. Boyle, L. E. Kidder, M. A. Scheel, B. Szilágyi, and N. W. Taylor, Suitability of hybrid gravitational waveforms for unequal-mass binaries, Phys. Rev. D **87**, 024009 (2013).

[69] P. Ajith *et al.*, The NINJA-2 catalog of hybrid post-Newtonian/numerical-relativity waveforms for non-precessing black-hole binaries, Classical Quantum Gravity **29**, 124001 (2012); Erratum **30**, 199401(E) (2013).

[70] M. Boyle, Uncertainty in hybrid gravitational waveforms: Optimizing initial orbital frequencies for binary black-hole simulations, Phys. Rev. D **84**, 064013 (2011).

[71] J. C. Bustillo, A. Boh, S. Husa, A. M. Sintes, M. Hannam *et al.*, Comparison of subdominant gravitational wave harmonics between post-Newtonian and numerical relativity calculations and construction of multi-mode hybrids, arXiv:1501.00918.

[72] M. Boyle, GWFrames, https://github.com/moble/GWFrames.

[73] L. Blanchet, T. Damour, G. Esposito-Farese, and B. R. Iyer, Gravitational Radiation from Inspiralling Compact Binaries Completed at the Third Post-Newtonian Order, Phys. Rev. Lett. **93**, 091101 (2004).

[74] L. Blanchet, Gravitational radiation from post-Newtonian sources and inspiralling compact binaries, Living Rev. Relativity **17**, 2 (2014).

[75] P. Jaranowskiand and G. Schfer, Dimensional regularization of local singularities in the 4th post-Newtonian two-point-mass Hamiltonian, Phys. Rev. D **87**, 081503 (2013).

[76] D. Biniand and T. Damour, Analytical determination of the two-body gravitational interaction potential at the 4th post-Newtonian approximation, Phys. Rev. D **87**, 121501 (2013).

[77] D. Biniand and T. Damour, High-order post-Newtonian contributions to the two-body gravitational interaction potential from analytical gravitational self-force calculations, Phys. Rev. D **89**, 064063 (2014).

[78] L. E. Kidder, Coalescing binary systems of compact objects to $(\text{post})^{5/2}$Newtonian order. V. Spin effects, Phys. Rev. D **52**, 821 (1995).

[79] C. M. Willand and A. G. Wiseman, Gravitational radiation from compact binary systems: Gravitational wave forms and energy loss to second post-Newtonian order, Phys. Rev. D **54**, 4813 (1996).

[80] A. Bohe, S. Marsat, G. Faye, and L. Blanchet, Next-to-next-to-leading order spin-orbit effects in the near-zone metric and precession equations of compact binaries, Classical Quantum Gravity **30**, 075017 (2013).

[81] M. Boyle, D. A. Brown, L. E. Kidder, A. H. Mroue, H. P. Pfeiffer, M. A. Scheel, G. B. Cook, and S. A. Teukolsky, High-accuracy comparison of numerical relativity simulations with post-Newtonian expansions, Phys. Rev. D **76**, 124038 (2007).

[82] L. Blanchet, G. Faye, B. R. Iyer, and S. Sinha, The third post-Newtonian gravitational wave polarisations and associated spherical harmonic modes for inspiralling compact binaries in quasi-circular orbits, Classical Quantum Gravity **25**, 165003 (2008).

[83] 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).

[84] G. Faye, L. Blanchet, and B. R. Iyer, Non-linear multipole interactions and gravitational-wave octupole modes for inspiralling compact binaries to third-and-a-half post-Newtonian order, Classical Quantum Gravity **32**, 045016 (2015).

[85] L. Blanchetand and G. Schäfer, Gravitational wave tails and binary star systems, Classical Quantum Gravity **10**, 2699 (1993).

[86] K. G. Arun, L. Blanchet, B. R. Iyer, and M. S. S. Qusailah, The 2.5PN gravitational wave polarizations from inspiralling compact binaries in circular orbits, Classical Quantum Gravity **21**, 3771 (2004).

[87] K. Barkett, Y. Chen, M. A. Scheel, and V. Varma, Detailed expansion of PN tidal splicing method (to be published).

[88] E. Jones, T. Oliphant, P. Peterson *et al.*, SciPy: Open source scientific tools for PYTHON (2001–), https://www.scipy.org/.

[89] S. Ossokine, M. Boyle, L. E. Kidder, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi, Comparing post-Newtonian and numerical relativity precession dynamics, Phys. Rev. D **92**, 104028 (2015).

[90] M. Favata, The gravitational-wave memory effect, Classical Quantum Gravity **27**, 084036 (2010).

[91] C. J. Handmer, B. Szilágyi, and J. Winicour, Gauge invariant spectral Cauchy characteristic extraction, Classical Quantum Gravity **32**, 235018 (2015).

[92] C. J. Handmer, B. Szilgyi, and J. Winicour, Spectral cauchy characteristic extraction of strain, news and gravitational radiation flux, Classical Quantum Gravity **33**, 225007 (2016).
[93] J. Winicour, Characteristic evolution and matching, Living Rev. Relativity **12**, 3 (2009).
[94] N. T. Bishop, R. Gómez, L. Lehner, and J. Winicour, Cauchy-characteristic extraction in numerical relativity, Phys. Rev. D **54**, 6153 (1996).
[95] M. Boyle, L. E. Kidder, S. Ossokine, and H. P. Pfeiffer, Gravitational-wave modes from precessing black-hole binaries, arXiv:1409.4431.
[96] P. Ajith, S. Babak, Y. Chen, M. Hewitson, B. Krishnan, A. M. Sintes, J. T. Whelan, B. Brügmann, P. Diener, N. Dorband, J. Gonzalez, M. Hannam, S. Husa, D. Pollney, L. Rezzolla, L. Santamaría, U. Sperhake, and J. Thornburg, A template bank for gravitational waveforms from coalescing binary black holes. I. Non-spinning binaries, Phys. Rev. D **77**, 104017 (2008).
[97] T. Damour, B. R. Iyer, and B. S. Sathyaprakash, A Comparison of search templates for gravitational waves from binary inspiral, Phys. Rev. D **63**, 044023 (2001); Erratum **72**, 029902(E) (2005).
[98] A. Buonanno, B. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash, Comparison of post-Newtonian templates for compact binary inspiral signals in gravitational-wave detectors, Phys. Rev. D **80**, 084043 (2009).
[99] S. E. Field, C. R. Galley, F. Herrmann, J. S. Hesthaven, E. Ochsner, and M. Tiglio, Reduced Basis Catalogs for Gravitational Wave Templates, Phys. Rev. Lett. **106**, 221102 (2011).
[100] F. Herrmann, S. E. Field, C. R. Galley, E. Ochsner, and M. Tiglio, Towards beating the curse of dimensionality for gravitational waves using reduced basis, Phys. Rev. D **86**, 084046 (2012).
[101] 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).
[102] 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).
[103] J. S. Hesthaven, B. Stamm, and S. Zhang, Efficient greedy algorithms for high-dimensional parameter spaces with applications to empirical interpolation and reduced basis methods, ESAIM: Math. Modell. Numer. Anal. **48**, 259 (2014).
[104] V. Varma, D. Gerosa, L. C. Stein, F. Hbert, and H. Zhang, High-Accuracy Mass, Spin, and Recoil Predictions of Generic Black-Hole Merger Remnants, Phys. Rev. Lett. **122**, 011101 (2019).
[105] P. Ajith, Addressing the spin question in gravitational-wave searches: Waveform templates for inspiralling compact binaries with nonprecessing spins, Phys. Rev. D **84**, 084037 (2011).
[106] C. Cutler and E. E. Flanagan, Gravitational waves from merging compact binaries: How accurately can one extract the binary's parameters from the inspiral wave form?, Phys. Rev. D **49**, 2658 (1994).
[107] E. Poissonand and C. M. Will, Gravitational waves from inspiraling compact binaries: Parameter estimation using second-post-Newtonian waveforms, Phys. Rev. D **52**, 848 (1995).
[108] D. J. A. McKechan, C. Robinson, and B. S. Sathyaprakash, A tapering window for time-domain templates and simulated signals in the detection of gravitational waves from coalescing compact binaries, Classical Quantum Gravity **27**, 084020 (2010).
[109] LIGO Scientific Collaboration, Advanced ligo anticipated sensitivity curves, https://dcc.ligo.org/LIGO-T0900288/public.
[110] B. Szilágyi, J. Blackman, A. Buonanno, A. Taracchini, H. P. Pfeiffer, M. A. Scheel, T. Chu, L. E. Kidder, and Y. Pan, Approaching the Post-Newtonian Regime with Numerical Relativity: A Compact-Object Binary Simulation Spanning 350 Gravitational-Wave Cycles, Phys. Rev. Lett. **115**, 031102 (2015).
[111] P. Kumar, K. Barkett, S. Bhagwat, N. Afshari, D. A. Brown, G. Lovelace, M. A. Scheel, and B. Szilgyi, Accuracy and precision of gravitational-wave models of inspiraling neutron star-black hole binaries with spin: Comparison with matter-free numerical relativity in the low-frequency regime, Phys. Rev. D **92**, 102001 (2015).
[112] T. Chu, H. Fong, P. Kumar, H. P. Pfeiffer, M. Boyle, D. A. Hemberger, L. E. Kidder, M. A. Scheel, and B. Szilágyi, On the accuracy and precision of numerical waveforms: effect of waveform extraction methodology, Classical Quantum Gravity **33**, 165001 (2016).
[113] M. A. Scheel, M. Giesler, D. A. Hemberger, G. Lovelace, K. Kuper, M. Boyle, B. Szilágyi, and L. E. Kidder, Improved methods for simulating nearly extremal binary black holes, Classical Quantum Gravity **32**, 105009 (2015).
[114] K. Chatziioannou, G. Lovelace, M. Boyle, M. Giesler, D. A. Hemberger, R. Katebi, L. E. Kidder, H. P. Pfeiffer, M. A. Scheel, and B. Szilgyi, Measuring the properties of nearly extremal black holes with gravitational waves, Phys. Rev. D **98**, 044028 (2018).
[115] S. A. Teukolsky, Perturbations of a rotating black hole. 1. Fundamental equations for gravitational electromagnetic and neutrino field perturbations, Astrophys. J. **185**, 635 (1973).
[116] S. A. Teukolsky, Rotating Black Holes: Separable Wave Equations for Gravitational and Electromagnetic Perturbations, Phys. Rev. Lett. **29**, 1114 (1972).
[117] E. Bertiand and A. Klein, Mixing of spherical and spheroidal modes in perturbed Kerr black holes, Phys. Rev. D **90**, 064012 (2014).
[118] E. Bertiand and A. Klein, http://www.phy.olemiss.edu/berti/qnms.html.