

Sparse Representations of Gravitational Waves from Precessing Compact Binaries

Jonathan Blackman,¹ Bela Szilagy, ¹ Chad R. Galley,¹ and Manuel Tiglio^{2,1}

¹*Theoretical Astrophysics, California Institute of Technology, Pasadena, California 91125, USA*

²*Center for Scientific Computation and Mathematical Modeling, and Joint Space Sciences Institute, Maryland Center for Fundamental Physics, Department of Physics, University of Maryland, College Park, Maryland 20742, USA*

(Received 29 January 2014; published 7 July 2014)

Many relevant applications in gravitational wave physics share a significant common problem: the seven-dimensional parameter space of gravitational waveforms from precessing compact binary inspirals and coalescences is large enough to prohibit covering the space of waveforms with sufficient density. We find that by using the reduced basis method together with a parametrization of waveforms based on their phase and precession, we can construct ultracompact yet high-accuracy representations of this large space. As a demonstration, we show that less than 100 judiciously chosen precessing inspiral waveforms are needed for 200 cycles, mass ratios from 1 to 10, and spin magnitudes ≤ 0.9 . In fact, using only the first 10 reduced basis waveforms yields a maximum mismatch of 0.016 over the whole range of considered parameters. We test whether the parameters selected from the inspiral regime result in an accurate reduced basis when including merger and ringdown; we find that this is indeed the case in the context of a nonprecessing effective-one-body model. This evidence suggests that as few as ~ 100 numerical simulations of binary black hole coalescences may accurately represent the seven-dimensional parameter space of precession waveforms for the considered ranges.

DOI: 10.1103/PhysRevLett.113.021101

PACS numbers: 04.25.dg, 04.25.Nx, 04.30.Db

Gravitational radiation produced by stellar-mass compact binaries (CB) are expected to be the main signals detected by the advanced generation of gravitational wave detectors [1–4]. Detecting these signals and estimating the parameters of their sources require sufficient sampling of the space of precessing CB waveforms. A CB intrinsically depends on its mass ratio and the spin angular momentum components of each body, which forms a seven-dimensional space [5].

Much progress has been made in sampling the subspace of nonspinning CB waveforms over the last decade. However, relevant applications such as building template banks for matched-filter searches, making parameter estimation studies, and modeling precessing CB by numerical simulations become prohibitively expensive in the full 7D space (e.g., see [6]). This phenomenon entails what is called the curse of dimensionality [7].

In this Letter, we show how to beat the curse of dimensionality for precessing CB inspirals. We find that only 50 judiciously chosen waveforms are needed to represent the 7D space with an accuracy from (9) better than 10^{-7} for 200 cycles, mass ratios $q \in [1, 10]$, dimensionless spin magnitudes $\|\vec{\chi}_{1,2}\| \leq 0.9$, and through $\ell = 8$ spherical harmonic modes. Using only the first 10 of these select waveforms yields a maximum representation error $\lesssim 1\%$. We find that for any given parameter range a remarkably small number of numerical relativity simulations of precessing binary black holes (BBH), if judiciously chosen, are sufficient to accurately represent any other precession waveform in that range. We expect these results

to be useful also for gravitational wave matched-filter searches and parameter estimation studies for CB coalescences.

Beating the curse of dimensionality.—Previous studies have shown that nonprecession subspaces of the full 7D waveforms space \mathcal{W} can be represented by linear spaces spanned by a relatively compact set of inspiral [8,9], ring-down [10], and inspiral-merger-ringdown (IMR) [11] waveforms, which form a reduced basis (RB). The RB waveforms are found by training a greedy algorithm [12,13] to learn from a given discretization of \mathcal{W} which are the most relevant waveforms for representing elements of \mathcal{W} with regard to a given error measure. See [8] for more details. The number of RB waveforms for nonprecessing inspirals hardly grows from two to four parameters [14]. Thus, one may beat the curse of dimensionality in the full 7D waveform space.

In this Letter, we construct a very compact but highly accurate RB of precession waveforms. We consider the following specifications on the 7D waveform space:

$$q = m_1/m_2 \in [1, 10], \quad \|\vec{\chi}_{1,2}\| \in [0, 0.9], \quad 200 \text{ cycles.} \quad (1)$$

These were chosen based partially on practical limitations of BBH simulations. However, the general message of this paper does not depend on our choice.

Key ingredients.—Our construction of a very sparse RB representation of precession waveforms depends on several key steps [15]: (i) A randomized resampling strategy [16] for training the greedy algorithm on the 7D waveform

space; (ii) a frame that rotates with the binary's precession; and (iii) a physically motivated parametrization of precession waveforms.

The first key ingredient is a modification of the standard greedy algorithm [8]. In its simplest inception, the greedy algorithm learns which waveforms can linearly span the space of interest in a nearly optimal way [12,13], starting from a sufficiently dense set of waveforms called a training set. However, the curse of dimensionality prevents us from sampling the waveform space with sufficient coverage to build a reliable training set. To overcome this, we randomly resample the 7D space from a uniform distribution using a fixed number K of waveforms at each iteration of the greedy algorithm. These waveforms constitute the training set at the current iteration. Because the 7D space is resampled at each iteration by different waveforms, the maximum error from projecting waveforms onto the current basis at the j th step is actually measuring this error over an effective training set with $j \times K$ randomly distributed waveforms. As more iterations are made, more of the 7D space is sampled and the more accurate the RB becomes. This is a simple implementation of more powerful techniques introduced in Ref. [16].

For our studies, we randomly and uniformly resampled $K \leq 36\,000$ waveforms at each iteration of the greedy algorithm. We began our studies with small K and increased each sample size up to $K = 36\,000$, for which the total number of RB waveforms was robust and independent of K . The largest training set used in our studies included more than 3×10^6 randomly selected waveforms.

The second key ingredient is to work in the binary's precessing frame instead of the usual inertial one. Specifically, we generate post-Newtonian (PN) precession waveforms in the time domain using the minimally rotating frame of Refs. [17–19]. In this frame, a precession waveform appears qualitatively similar to waveforms from nonprecessing binaries in their inertial frame [6,17,19–23], resulting in a weaker dependence on parameters than they have in the inertial frame. The rotation from the minimally rotating frame to the inertial one and vice versa can be accounted for by any convenient representation of the $SO(3)$ group.

The third key ingredient, and perhaps the most crucial, is that we choose to parametrize precession waveforms by their phase instead of by time or frequency. To motivate this choice we momentarily consider the frequency-domain waveform (in the stationary phase approximation) for a nonspinning binary inspiral at leading order (“OPN”) in the PN approximation,

$$h(f; \mathcal{M}) = A\mathcal{M}^{5/6} f^{-7/6} e^{i\Phi_0(f; \mathcal{M})}, \quad (2)$$

where $\mathcal{M} = M\nu^{3/5}$ is the chirp mass, M is the total mass, ν is the symmetric mass ratio, A is a constant independent of the binary's intrinsic parameters, and

$$\Phi_0(f; \mathcal{M}) \equiv \frac{3}{128} (\pi\mathcal{M}f)^{-5/3}. \quad (3)$$

Reparametrizing (2) by its phase, now taken as the independent variable, gives

$$H(\varphi; \mathcal{M}) \equiv h(F(\varphi); \mathcal{M}) = A'\mathcal{M}^2\varphi^{7/10} e^{i\varphi}, \quad (4)$$

with $A' = A\pi^{7/6}(128/3)^{7/10}$ and $F(\varphi)$ from solving $\Phi_0(f = F) = \varphi$. In this phase domain, all waveforms are proportional to each other, thus constituting a 1D space. In fact, performing the greedy algorithm analytically (versus numerically) returns a single basis waveform that *exactly* represents all such waveforms in the continuum. This is the intrinsic dimensionality of the problem as has long been known because OPN waveforms only depend on the chirp mass. Therefore, a single RB waveform spans the whole OPN waveform space. To close the system, we also need to represent the mapping between the phase and frequency domains,

$$F(\varphi; \mathcal{M}) = \frac{1}{\pi\mathcal{M}} \left(\frac{128\varphi}{3} \right)^{-3/5}, \quad (5)$$

using a separate basis. As we see again, the frequencies for different chirp masses are all proportional to each other. Therefore, any OPN waveform, as a function of frequency, is represented by *one* RB waveform through the nonlinear transformation in (3).

For comparison's sake, we implemented a standard greedy algorithm following [8] using OPN waveforms parametrized by frequency (not phase) for binaries with a fixed total mass and with mass ratios and number of cycles as in (1). We found that 152 RB waveforms are required to reach numerical round-off errors in representing any waveform in this 1D space and 138 RB waveforms to reach 1% errors. Using the phase parametrization yields a single RB waveform for exact representation whereas a frequency parametrization can give a much larger RB for *approximate* representation.

Part of the reason why using waveforms in the phase domain (or φ domain) is advantageous is because the oscillations in two waveforms *always* cancel in the scalar product used to measure the projection error onto the RB in the greedy algorithm,

$$\langle H_{\mathcal{M}_1}, H_{\mathcal{M}_2} \rangle_\varphi \equiv \int_{\varphi_{\min}}^{\varphi_{\max}} d\varphi H(\varphi; \mathcal{M}_1) H^*(\varphi; \mathcal{M}_2). \quad (6)$$

For OPN waveforms this results in a very smooth dependence on the chirp masses since (6) is $\propto \mathcal{M}_1^2 \mathcal{M}_2^2$. Similarly, the waveform frequency as a function of phase (5) has a very smooth dependence on them as well.

Higher PN orders include more physics, such as the nonlinearity of general relativity and spin orbit, spin-1–spin-2, and self-spin interactions, that depend on all 7

intrinsic parameters. These contributions add more structure to the waveforms but only weakly depend on the parameters. This is especially true in the φ domain and, as discussed below, we also find this holds through the merger and ringdown phases where the PN expansion parameter is no longer small. Consequently, one may expect to find only a relatively small number \mathcal{N} of RB waveforms, possibly as few as $\mathcal{N} = O(d)$.

Method outline.—In this Letter, we use 3.5 PN precessing inspiral waveforms. We solve the PN equations (see Ref. [24] and references therein) using the approach of Refs. [18,19] where the waveforms themselves are solved in a frame that minimizes the binary’s precession, along with a rotation operator represented by unit quaternions to track this frame relative to the fiducial inertial frame. All waveforms in this minimally rotating frame are normalized to unity, and the initial orbital phases are aligned. It is convenient to decompose the waveform into spin-weighted spherical harmonic modes [25] characterized by (ℓ, m) . We use the phase associated with the $(\ell, m) = (2, 2)$ mode to parametrize the waveform but other choices are possible. A precession waveform $h(t)$ in the inertial frame is thus decomposed in the following way:

$$h(t) \rightarrow (\{H_{\ell m}(\varphi)\}, T(\varphi), Q(\varphi)), \quad (7)$$

where $H_{\ell m}$ is a spin-weighted spherical harmonic mode in the minimally rotating frame, $T(\varphi)$ is the function relating the (2,2) phase to the time coordinate, and Q is the unit quaternion describing the rotation back to the inertial frame. We take into account all modes up to $\ell = 8$ and cut all waveforms off at a dimensionless frequency of 0.2 in the (2,2) modes. Finally, all waveforms contain 200 wave cycles.

We build a RB for each component in the decomposition (7). It is natural to use the scalar product in (6) for the T and Q functions but to integrate the minimally rotating waveform over all angles or orientations on the two sphere so that, upon using the orthogonality of the spin-weighted spherical harmonics,

$$\langle H_{\lambda_1}, H_{\lambda_2} \rangle_{\varphi} \equiv \sum_{\ell, m} \int_{\varphi_{\min}}^{\varphi_{\max}} d\varphi H_{\ell m}(\varphi; \lambda_1) H_{\ell m}^*(\varphi; \lambda_2), \quad (8)$$

where λ_i is a tuple of parameter values. Executing a greedy algorithm on each component in (7) would result in a selection of parameter values that are different for each element. To choose the same parameters for all three RB, we define a total projection error ϵ_{φ} through

$$\epsilon_{\varphi}(\lambda) \equiv 8 \times 10^{-6} \|\delta T_{\lambda}\|_{\varphi}^2 + 0.5 \|\delta H_{\lambda}\|_{\varphi}^2 + 0.0031 \|\delta Q_{\lambda}\|_{\varphi}^2, \quad (9)$$

so as to receive approximately equal contributions from each component. Here, $\lambda = (q, \vec{\chi}_1, \vec{\chi}_2)$ is a tuple of 7D parameter values, $\delta X_{\lambda} = X_{\lambda} - P_X[X_{\lambda}]$ with X one of the

elements in (7), and P_X the projection operator onto the basis for X . The numerical coefficients are fixed to give approximately equal contributions to the mismatch in the time domain and inertial frame in the case of small random perturbations. Binaries with periods near $200M$ lead to a small coefficient for the time function.

Results for precessing binary inspirals.—We implemented a greedy algorithm using the three key ingredients discussed above to find RB representations for the space of precession waveforms as in (1). The left panel of Fig. 1 shows the maximum of the total projection error (9) found at each iteration of the greedy algorithm. We observe a power-law decay with exponent ≈ -8 . The total error is not monotonically decreasing because of the constant resampling at each iteration. We observe that the maximum normalized projection error over the training set is 10^{-2} using 10 basis waveforms and $\approx 4 \times 10^{-8}$ for 50. Also shown are the contributions to the total error from the projections onto the basis of each component in (7).

To measure the error in the time-domain inertial frame between a waveform h and its RB approximation h_{app} we use the standard mismatch:

$$\text{Mismatch} = 1 - \text{Re} \langle h, h_{\text{app}} \rangle_t, \quad (10)$$

where for the two functions A, B the time-domain complex scalar product is $\langle A, B \rangle_t \equiv \int_{t_{\min}}^{t_{\max}} dt A(t) B^*(t)$. In order to

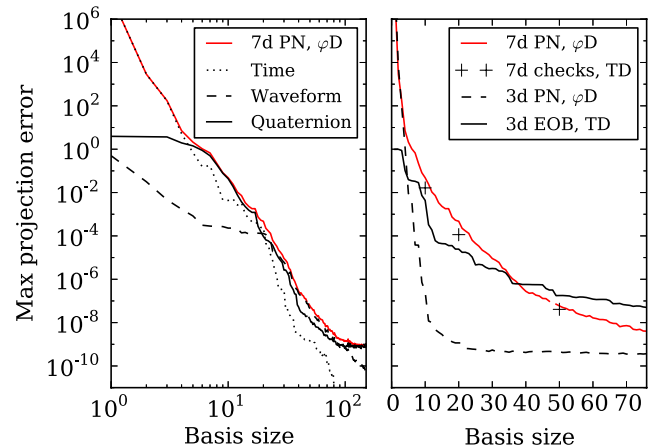


FIG. 1 (color online). Left: maximum φ -domain (φD) projection error (red) from (9) for 7D post-Newtonian precession waveforms versus basis size. The contributions from the time function (dotted), waveform in the minimally rotating frame (dashed), and quaternion (solid) are also shown. Right: maximum time-domain, inertial frame mismatches from (10) for 10^7 randomly selected waveforms (+) using the first 10, 20, and 50 RB waveforms. Also plotted are φ -domain projection errors for non-precessing PN waveforms (dashed) and the time-domain (TD) projection errors from using the latter parameter values selected by the greedy algorithm to represent EOB waveforms (solid black), which additionally include merger and ringdown phases.

measure the quality of the RB approximation itself, we do not extremize the mismatch with respect to the relative phase and time shift between h and h_{app} .

It is not obvious that the basis generated using (9) from the minimally rotating frame and φ domain will be accurate for inertial frame waveforms expressed in the time domain. We find that the φ domain, precessing basis is highly accurate for representing time-domain, inertial-frame waveforms. The right panel in Fig. 1 shows the mismatch (+) from using the first 10, 20, and 50 basis functions to represent more than 10^7 randomly chosen waveforms as in (1). The corresponding maximum mismatches are 0.016, 1.1×10^{-4} , and 4.1×10^{-8} , respectively.

The table below shows that, for a given error, the number of RB waveforms needed to accurately represent the subspace of \mathcal{W} with the indicated dimension d grows approximately linearly with d , not exponentially. The first three dimensions are the mass ratio q (1D), with the z components of the first (2D) then second spins (3D) included.

Error	Basis size			
	1D	2D	3D	7D
$\lesssim 10^{-2}$	4	6	7	13
$\lesssim 10^{-4}$	4	7	8	20
$\lesssim 3 \times 10^{-8}$	6	15	23	50

Figure 2 shows the first 90 parameters selected by our greedy algorithm and presented according to which component—time, minimally rotating waveform, quaternion—is the dominant contribution to the total representation error in the left panel of Fig. 1. The spins' components are taken at the initial time where the inertial and minimally rotating frames are equal. The mass ratios dominate the endpoints of the considered interval in (1). Both spins' magnitudes tend to be in $[0.8, 0.9]$. The projections of the spins onto the initial orbital angular momentum seem to be anticorrelated, at least when the waveform contribution to (9) is dominant. The x - y components of the spins tend to lie on a circle for the smaller mass m_2 ; there is less clear structure for m_1 .

From inspiral to coalescence.—We test whether the parameters selected from the inspiral regime result in an accurate RB when including merger and ringdown. This issue has immediate relevance for building a RB for expensive numerical relativity simulations of precessing BBH mergers that, in turn, has important ramifications for data analysis applications with gravitational wave detectors and for modeling merger simulations [11].

Currently, we can answer the above question for spinning but nonprecessing BBH coalescences, which involves only the three parameters q , $|\vec{\chi}_1|$, and $|\vec{\chi}_2|$, for which an effective-one-body (EOB) semianalytical model [26] of IMR is available [27,28]. We first used our greedy

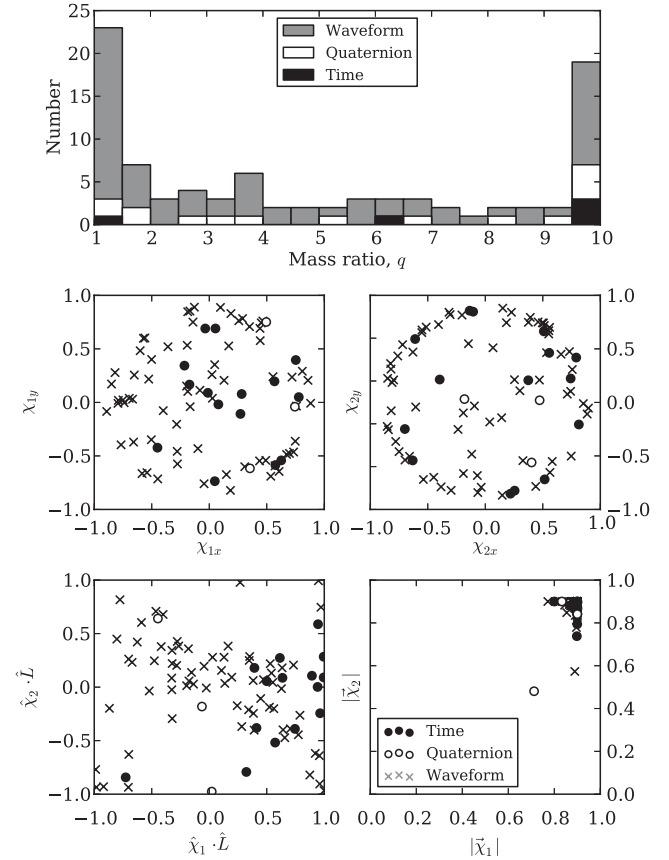


FIG. 2. Mass ratios (top), x - y components of both spins (second row), projection of inertial frame spins onto initial orbital angular momentum unit vector (bottom left), and both spin magnitudes (bottom right) selected by our greedy algorithm.

algorithm to find the parameters for building a RB for the nonprecessing inspiral PN waveforms using the φ domain error in (9). We then generated a basis using the EOB nonprecessing coalescence waveforms evaluated at those selected parameters. Last, we randomly generated more than $10^{d=3}$ EOB waveforms and computed the time-domain inertial frame mismatch from (10). The results of this study are shown as the solid black curve in the right panel of Fig. 1. For the first 20 inspiral RB waveforms, the maximum mismatch of the EOB waveforms is about 3×10^{-5} while for the first 50 it is about 2×10^{-7} .

Outlook.—Based on traditional methods to sample the waveform space, which scale exponentially with dimension [29–32], it has been perceived that an intractable number of numerical relativity simulations would be needed to represent the space of BBH for any given number of orbits. We have found evidence that a remarkably small number of numerical relativity BBH simulations may actually be needed, if judiciously chosen, to build a high accuracy RB to represent the whole space of interest.

Based on the nonprecessing EOB results presented above, performing numerical simulations of BBH mergers for the first 50–90 parameters selected by our greedy

algorithm may be sufficient to represent the precession waveforms of any other coalescences in the parameter ranges of (1). This constitutes less than one tenth of the number of randomly chosen simulations performed over the last few years [22,33–35]. This work suggests that an unexpectedly small number of low-mass inspiral waveforms may represent the frequency and parameter ranges of interest to gravitational wave detectors, which may also enable very compact reduced-order quadratures [36,37] of overlap integrals for fast parameter estimation studies. Finally, this work opens the door for building surrogate models [11] of precessing inspiral waveforms that can be useful for multiple query applications in place of solving a large number of parametrized ordinary differential equations.

We thank Scott Field and Rory Smith for comments and discussions. We thank Michael Boyle for permission to use his Triton code for solving precessing PN waveforms, now replaced by his public open source code available from [38]. We have also used the LAL (LSC Algorithm Library) EOB code version 6.11.0.1 [39]. This project was supported in part by the Fairchild Foundation, NSF Grants No. PHY-1068881, No. CAREER PHY-0956189, and No. PHY-1005655 to Caltech, NASA Grant No. NNX10AC69G, and NSF Grants No. PHY-1208861, No. PHY-1316424, and No. PHY-1005632 to the University of Maryland. Computations were performed on the Zwicky cluster at Caltech, which is supported by the Sherman Fairchild Foundation and by NSF Grant No. PHY-0960291. This work used NSF XSEDE resources under allocation No. TG-PHY990007N.

-
- [1] G. M. Harry *et al.* (LIGO Scientific Collaboration), *Classical Quantum Gravity* **27**, 084006 (2010).
- [2] G. Losurdo *et al.* (LIGO Scientific Collaboration and VIRGO Collaboration), *J. Phys. Conf. Ser.* **110**, 062016 (2008).
- [3] K. Somiya *et al.* (KAGRA Collaboration), *Classical Quantum Gravity* **29**, 124007 (2012).
- [4] INDIGO, <http://www.gw-indigo.org>.
- [5] Gravitational wave detectors have a finite frequency bandwidth that introduces a total mass scale, thus adding an 8th parameter that we will ignore. We also focus on quasicircular inspirals.
- [6] M. Hannam, [arXiv:1312.3641](https://arxiv.org/abs/1312.3641) [*Gen. Relativ. Gravit.* (to be published)].
- [7] R. Bellman, *Adaptive Control Process: A Guided Tour* (Princeton University Press, Princeton, NJ, 1961).
- [8] S. E. Field, C. R. Galley, F. Herrmann, J. S. Hesthaven, E. Ochsner, and M. Tiglio, *Phys. Rev. Lett.* **106**, 221102 (2011).
- [9] S. E. Field, C. R. Galley, and E. Ochsner, *Phys. Rev. D* **86**, 084046 (2012).
- [10] S. Caudill, S. E. Field, C. R. Galley, F. Herrmann, and M. Tiglio, *Classical Quantum Gravity* **29**, 095016 (2012).
- [11] S. E. Field, C. R. Galley, J. S. Hesthaven, J. Kaye, and M. Tiglio, [arXiv:1308.3565](https://arxiv.org/abs/1308.3565) [*Phys. Rev. X* (to be published)].
- [12] P. Binev, A. Cohen, W. Dahmen, R. DeVore, G. Petrova, and P. Wojtaszczyk, *SIAM J. Math. Anal.* **43**, 1457 (2011).
- [13] R. DeVore, G. Petrova, and P. Wojtaszczyk, *Constr. Approx.* **37**, 455 (2013).
- [14] C. R. Galley, F. Herrmann, J. Silberholz, M. Tiglio, and G. Guerberoff, *Classical Quantum Gravity* **27**, 245007 (2010).
- [15] The last two ingredients can be viewed as aspects of nonlinear dimensional reduction and manifold learning, which aim to reveal the intrinsic dimensionality of large amounts of data.
- [16] J. S. Hesthaven, B. Stamm, and S. Zhang, *Math. Modell. Numer. Anal.* **48**, 259 (2014).
- [17] P. Schmidt, M. Hannam, S. Husa, and P. Ajith, *Phys. Rev. D* **84**, 024046 (2011).
- [18] M. Boyle, R. Owen, and H. P. Pfeiffer, *Phys. Rev. D* **84**, 124011 (2011).
- [19] M. Boyle, *Phys. Rev. D* **87**, 104006 (2013).
- [20] P. Schmidt, M. Hannam, and S. Husa, *Phys. Rev. D* **86**, 104063 (2012).
- [21] A. Lundgren and R. O’Shaughnessy, *Phys. Rev. D* **89**, 044021 (2014).
- [22] L. Pekowsky, R. O’Shaughnessy, J. Healy, and D. Shoemaker, *Phys. Rev. D* **88**, 024040 (2013).
- [23] M. Hannam *et al.*, [arXiv:1308.3271](https://arxiv.org/abs/1308.3271).
- [24] A. Bohe, S. Marsat, G. Faye, and L. Blanchet, *Classical Quantum Gravity* **30**, 075017 (2013).
- [25] K. Thorne, *Rev. Mod. Phys.* **52**, 299 (1980).
- [26] This EOB model includes only the quadrupole mode, which is dominant through most of the IMR but less so near merger.
- [27] A. Buonanno and T. Damour, *Phys. Rev. D* **59**, 084006 (1999).
- [28] A. Taracchini, Y. Pan, A. Buonanno, E. Barausse, M. Boyle, T. Chu, G. Lovelace, H. P. Pfeiffer, and M. A. Scheel, *Phys. Rev. D* **86**, 024011 (2012).
- [29] B. J. Owen, *Phys. Rev. D* **53**, 6749 (1996).
- [30] B. J. Owen and B. S. Sathyaprakash, *Phys. Rev. D* **60**, 022002 (1999).
- [31] I. W. Harry, B. Allen, and B. S. Sathyaprakash, *Phys. Rev. D* **80**, 104014 (2009).
- [32] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, and A. Weinstein, *Phys. Rev. D* **89**, 084041 (2014).
- [33] P. Ajith *et al.*, *Classical Quantum Gravity* **29**, 124001 (2012).
- [34] I. Hinder *et al.*, *Classical Quantum Gravity* **31**, 025012 (2014).
- [35] A. H. Mroue *et al.*, *Phys. Rev. Lett.* **111**, 241104 (2013).
- [36] H. Antil, S. E. Field, F. Herrmann, R. H. Nochetto, and M. Tiglio, *J. Sci. Comput.* **57**, 604 (2013).
- [37] P. Canizares, S. E. Field, J. R. Gair, and M. Tiglio, *Phys. Rev. D* **87**, 124005 (2013).
- [38] GITHUB, <https://github.com/MOBlE/PostNewtonian>.
- [39] LALSUITE, <https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html>.