

Gravitational wave peak luminosity model for precessing binary black holes

Afura Taylor^{1,*} and Vijay Varma^{2,1,†}

¹*TAPIR 350-17, California Institute of Technology, 1200 East California Boulevard, Pasadena, California 91125, USA*

²*Department of Physics, and Cornell Center for Astrophysics and Planetary Science, Cornell University, Ithaca, New York 14853, USA*



(Received 1 October 2020; accepted 27 October 2020; published 18 November 2020)

When two black holes merge, a tremendous amount of energy is released in the form of gravitational radiation in a short span of time, making such events among the most luminous phenomenon in the Universe. Models that predict the peak luminosity of black hole mergers are of interest to the gravitational wave community, with potential applications in tests of general relativity. We present a surrogate model for the peak luminosity that is directly trained on numerical relativity simulations of precessing binary black holes. Using Gaussian process regression, we interpolate the peak luminosity in the seven-dimensional parameter space of precessing binaries with mass ratios $q \leq 4$ and spin magnitudes $\chi_1, \chi_2 \leq 0.8$. We demonstrate that our errors in estimating the peak luminosity are lower than those of existing fitting formulas by about an order of magnitude. In addition, we construct a model for the peak luminosity of aligned-spin binaries with mass ratios $q \leq 8$ and spin magnitudes $|\chi_{1z}|, |\chi_{2z}| \leq 0.8$. We apply our precessing model to infer the peak luminosity of the GW event GW190521 and find the results to be consistent with previous predictions.

DOI: [10.1103/PhysRevD.102.104047](https://doi.org/10.1103/PhysRevD.102.104047)

I. INTRODUCTION

As the gravitational wave (GW) detectors LIGO [1] and Virgo [2] approach their design sensitivity, GW detections are becoming routine [3–7]. Binary black hole (BBH) mergers are the most abundant source for these detectors. Such mergers provide a unique laboratory for studying black hole (BH) astrophysics as well as for testing general relativity. At the time of merger, the BHs are moving at about half the speed of light, and the spacetime is highly dynamical. As a result, for a brief moment, BBH mergers are among the most luminous events in the Universe. For example, the recently announced GW event GW190521 [7] radiated approximately $7.6 M_\odot$ of energy in GWs in a fraction of a second, reaching a peak luminosity of approximately $208 M_\odot c^2/s = 3.7 \times 10^{56}$ erg/s [8].

The above estimate is obtained by applying peak luminosity models [9,10] based on numerical relativity (NR) simulations to the measured masses and spins of the component BHs. Apart from predicting the peak luminosity of GW events, such models can be used to understand the impact of supermassive BH mergers on circumbinary accretion disks [11] and possible electromagnetic counterparts [12,13]. In addition, one can test general relativity by independently estimating the peak luminosity through a

theory-independent signal reconstruction [14,15] and comparing with the prediction from NR. A similar test was performed for the peak frequency in Ref. [16]. As detector sensitivity improves, these applications will need accurate models that capture the full physics of the NR simulations.

NR simulations are essential to model the BH merger process and the resulting GW peak luminosity. However, these are prohibitively expensive for most GW data analysis applications. As a result, various phenomenological fits have been developed for the peak luminosity [9,10,17,18]; starting with an ansatz, these models calibrate any free coefficients to NR simulations. However, all of these models are restricted to aligned-spin systems, where the BH spins are aligned to the orbital angular momentum direction (\hat{L}). For generic binaries, however, the spins can be tilted with respect to \hat{L} . For these systems, the spins interact with the orbit (and each other), leading to precession of the orbital plane and the spins [19]. Precession causes modulations in the GW signal, and as a result, the peak luminosity is affected.

In this paper, we present a Gaussian process regression (GPR) based NR surrogate model for the peak luminosity of generically precessing BBHs. NR surrogate models directly interpolate between NR simulations rather than assume an ansatz about the underlying phenomenology. These methods have been successfully used to model the GW signal [20–22] as well as the remnant BH properties [20,23,24] of precessing BBHs. Through cross-validation

*afura.taylor@gmail.com

†vvarma@cornell.edu

studies, these models have been shown to approach the accuracy level of the NR simulations themselves.

In particular, we present two models:

- (1) NRSur7dq4Remnant: a seven-dimensional precessing model trained against systems with mass ratios $q \leq 4$,¹ dimensionless spin magnitudes $\chi_1, \chi_2 \leq 0.8$, and generic spins directions.
- (2) NRSur3dq8Remnant: a three-dimensional aligned-spin model trained against systems with mass ratios up to $q \leq 8$ and aligned spins $|\chi_{1z}|, |\chi_{2z}| \leq 0.8$.

We use the same names, respectively, as the precessing remnant model of Ref. [20] and the aligned-spin remnant model of Ref. [23], as we make the models available in the same interface through the publicly available PYTHON module SURFINBH [25]. Even though peak luminosity is not technically a property of the remnant BH, we expect that using the same interface will make using the models easier for our users.

The rest of the paper is as follows. We describe our fitting procedure in Sec. II. In Sec. III, we compare the models against NR simulations to assess their accuracy. In Sec. IV, we apply our precessing model to predict the peak luminosity of GW190521. We end with some concluding remarks in Sec. V.

II. MODELING METHODS

The GW luminosity is defined as [18]

$$\mathcal{L}(t) = \frac{1}{16\pi} \sum_{\ell, m} \left| \lim_{r \rightarrow \infty} (r \dot{\hbar}_{\ell m}) \right|^2, \quad (1)$$

where the dot represents a time derivative, the $\|$ represents the absolute value, and $\hbar_{\ell m}$ represents the complex the spin-2 weighted spherical harmonic mode with indices (ℓ, m) . We use the time derivative of $r \hbar_{\ell m}$ extrapolated to future null infinity [26] in the place of $\lim_{r \rightarrow \infty} (r \dot{\hbar}_{\ell m})$. The extrapolated strain data are obtained from NR simulations performed with the Spectral Einstein Code (SPEC) [27] code, available through the Simulating Extreme Spacetimes (SXS) [28] Catalog [29,30]. The strain data are first interpolated onto a uniform time array (with step size $0.1M$, where M is the total mass) using cubic splines. Then, we use a fourth-order finite-difference derivative to get the time derivative of the strain.

We determine the peak luminosity as

$$\mathcal{L}_{\text{peak}} = \max_t \mathcal{L}(t), \quad (2)$$

where we determine the peak value by fitting a quadratic function to five adjacent samples of $\mathcal{L}(t)$, consisting of the largest sample and two neighbors on either side. Before

¹We use the convention $q = m_1/m_2 \geq 1$, where m_1 (m_2) is the mass of the heavier (lighter) BH.

applying our fitting method, we first take a logarithm of the peak luminosity and model $\log(\mathcal{L}_{\text{peak}})$. We find that this leads to more accurate fits than directly modeling $\mathcal{L}_{\text{peak}}$. When the model is evaluated, we can easily get the predicted peak luminosity by taking the exponential of the fit output.

For the aligned-spin model NRSur3dq8Remnant, we include the $\ell \leq 4$ and (5,5) modes but not the (4,1) or (4,0) modes in Eq. (1). We include the $m > 0$ modes twice to account for the $m < 0$ modes, which are given by $\hbar_{\ell, -m} = (-1)^\ell \hbar_{\ell m}^*$ due to the symmetries of aligned-spin systems. The included modes are the same as those used for the surrogate model of Ref. [31]. The reason for excluding the (4,1), (4,0), ($\ell = 5, m < 5$), and $\ell > 5$ modes is twofold:

- (i) These modes have very small amplitudes and do not contribute significantly to the sum in Eq. (1).
- (ii) The small amplitude of some of these modes [particularly (4,1) and (4,0)] can behave poorly when extrapolated [26].

We expect that this will be resolved in the future with Cauchy characteristic extraction [32,33].

For the precessing model NRSur7dq4Remnant, we use all $\ell \leq 5$ modes. Due to the orbital precession, even modes like (4,1), (4,0), and ($\ell = 5, m < 5$) can have significant amplitude due to mode mixing (see, e.g., Ref. [20]). Therefore, these modes behave reasonably well when extrapolated. Note that the $m < 0$ modes are directly included when doing the sum in Eq. (1) as the aforementioned symmetry for $m < 0$ does not hold for precessing systems.

A. Gaussian process regression

We construct fits in this work using GPR [34] as implemented in SCIKIT-LEARN [35]. We closely follow the procedure outlined in the supplement of Ref. [23], which we describe briefly in the following.

We start with a training set of n observations, $\mathcal{TS} = \{\Lambda^i, f(\Lambda^i) | i = 1, \dots, n\}$, where each Λ^i denotes an input vector of dimension D and $f(\Lambda^i)$ is the corresponding scalar output. In our case, Λ is given by Eqs. (6) and (11), respectively, for the precessing and aligned-spin models, and $f(\Lambda) = \log(\mathcal{L}_{\text{peak}})$. Our goal is to use \mathcal{TS} to make predictions for the underlying $f(\Lambda)$ at any point Λ_* that is not in \mathcal{TS} .

A Gaussian process (GP) can be thought of as a probability distribution of functions. More formally, a GP is a collection of random variables, any finite number of which have a joint Gaussian distribution [34]. A GP is completely specified by its mean function $m(\Lambda)$ and covariance function $k(\Lambda, \Lambda')$, i.e., $f(\Lambda) \sim \mathcal{GP}(m(\Lambda), k(\Lambda, \Lambda'))$. Consider a prediction set of n_* test inputs and their corresponding outputs (which are unknown): $\mathcal{PS} = \{\Lambda_*^i, f(\Lambda_*^i) | i = 1, \dots, n_*\}$. By the definition of a GP, outputs of \mathcal{TS} and \mathcal{PS}

[respectively, $\mathbf{f} = \{f(\Lambda^i)\}$, $\mathbf{f}_* = \{f(\Lambda_*^i)\}$] are related by a joint Gaussian distribution,

$$\begin{bmatrix} \mathbf{f} \\ \mathbf{f}_* \end{bmatrix} = \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} K_{\Lambda\Lambda} & K_{\Lambda\Lambda_*} \\ K_{\Lambda_*\Lambda} & K_{\Lambda_*\Lambda_*} \end{bmatrix}\right), \quad (3)$$

where $K_{\Lambda\Lambda_*}$ denotes the $n \times n_*$ matrix of the covariance $k(\Lambda, \Lambda_*)$ evaluated at all pairs of training and prediction points, and similarly for the other K matrices.

Equation (3) provides the Bayesian prior distribution for \mathbf{f}_* . The posterior distribution is obtained by restricting this joint prior to contain only those functions which agree with the observed data points [34], i.e.,

$$p(\mathbf{f}_* | \mathcal{TS}) = \mathcal{N}(K_{\Lambda_*\Lambda} K_{\Lambda\Lambda}^{-1} \mathbf{f}, K_{\Lambda_*\Lambda_*} - K_{\Lambda_*\Lambda} K_{\Lambda\Lambda}^{-1} K_{\Lambda\Lambda_*}). \quad (4)$$

The mean of this posterior provides an estimator for $f(\Lambda)$ at Λ_* , while its width is the prediction error.

Finally, one needs to specify the covariance (or kernel) function $k(\Lambda, \Lambda')$. Following Ref. [23], we implement the following kernel,

$$k(\Lambda, \Lambda') = \sigma_k^2 \exp\left[-\sum_{j=1}^D \left(\frac{\Lambda^j - \Lambda'^j}{\sqrt{2}\sigma_j}\right)^2\right] + \sigma_n^2 \delta_{\Lambda\Lambda'}, \quad (5)$$

where $\delta_{\Lambda\Lambda'}$ is the Kronecker delta. In words, we use a product between a squared exponential kernel (parametrized by σ_j) and a constant kernel (parametrized by σ_k^2), to which we add a white kernel (parametrized by σ_n^2) to account for additional noise in the training data [34,35].

GPR fit construction involves determining the $D+2$ hyperparameters (σ_k , σ_n , and σ_j) which maximize the marginal likelihood of the training data under the GP prior [34]. Local maxima are avoided by repeating the optimization with ten different initial guesses, obtained by sampling uniformly in log in the hyperparameter space described below.

Before constructing the GPR fit, we preprocess the training data as follows. We first subtract a linear fit and the mean of the resulting values. The data are then normalized by dividing by the standard deviation of the resulting values. The inverse of these transformations is applied at the time of the fit evaluation. The reasoning behind the preprocessing is twofold:

- (1) The demeaning and normalization allows us to apply the same ranges (described below) for the GPR hyperparameters for a wide range of models. For instance, we used the same settings to model the remnant BH properties in Refs. [20,23].
- (2) The data are simpler to model after removing the linear component, leading to more accurate fits.

For each dimension of Λ , we define $\Delta\Lambda^j$ to be the range of the values of Λ^j in \mathcal{TS} and consider $\sigma_j \in [0.01 \times \Delta\Lambda^j, 10 \times \Delta\Lambda^j]$. Larger length scales are unlikely to be relevant

and smaller length scales are unlikely to be resolvable. The remaining hyperparameters are sampled in $\sigma_k^2 \in [10^{-2}, 10^2]$ and $\sigma_n^2 \in [10^{-7}, 10^{-2}]$. These choices are meant to be conservative and are based on prior exploration of the typical magnitude and noise level in our training data.

B. Precessing model, NRSur7dq4Remnant

For precessing systems the parameter space is seven dimensional comprising the mass q and two spin 3-vectors χ_1 and χ_2 . Here, $q = m_1/m_2$ is the mass ratio with $m_1 \geq m_2$, and χ_1 (χ_2) is the dimensionless spin vector of the heavier (lighter) BH. The total mass ($M = m_1 + m_2$) scales out of the problem and does not constitute an additional parameter for modeling. We use the 1528 NR waveforms used for the surrogate models of Ref. [20], which cover the parameter space $q \leq 4, \chi_1, \chi_2 \leq 0.8$, where χ_1 (χ_2) is the magnitude of χ_1 (χ_2).

Following Refs. [20,23], we parametrize the precessing fit using the coorbital frame spins $\chi_{1,2}^{\text{coorb}}$ at $t = -100M$ before the peak of the total waveform amplitude (as defined in Eq. 5 of Ref. [20]). The coorbital frame is a time-dependent noninertial frame in which the z axis is along the instantaneous \hat{L} direction and x axis is along the instantaneous line of separation between the BHs, with the heavier BH on the positive x axis.² The NRSur7dq4Remnant fit is parametrized as follows,

$$\Lambda = [\log(q), \chi_{1x}^{\text{coorb}}, \chi_{1y}^{\text{coorb}}, \hat{\chi}^{\text{coorb}}, \chi_{2x}^{\text{coorb}}, \chi_{2y}^{\text{coorb}}, \chi_a^{\text{coorb}}], \quad (6)$$

where $\hat{\chi}^{\text{coorb}}$ is the spin parameter entering the GW phase at leading order [36–39] in the post-Newtonian (PN) expansion

$$\hat{\chi}^{\text{coorb}} = \frac{\chi_{\text{eff}}^{\text{coorb}} - 38\eta(\chi_{1z}^{\text{coorb}} + \chi_{2z}^{\text{coorb}})/113}{1 - 76\eta/113}, \quad (7)$$

$$\chi_{\text{eff}}^{\text{coorb}} = \frac{q\chi_{1z}^{\text{coorb}} + \chi_{2z}^{\text{coorb}}}{1 + q}, \quad (8)$$

$$\eta = \frac{q}{(1 + q)^2}, \quad (9)$$

and χ_a^{coorb} is the ‘‘antisymmetric spin,’’

$$\chi_a^{\text{coorb}} = \frac{1}{2}(\chi_{1z}^{\text{coorb}} - \chi_{2z}^{\text{coorb}}). \quad (10)$$

We empirically found this parametrization to perform more accurately than the more intuitive choice $\hat{\Lambda} = [q, \chi_{1x}^{\text{coorb}}, \chi_{1y}^{\text{coorb}}, \chi_{1z}^{\text{coorb}}, \chi_{2x}^{\text{coorb}}, \chi_{2y}^{\text{coorb}}, \chi_{2z}^{\text{coorb}}]$ used in Ref. [21].

²Here, the BH positions are defined using the waveform at future null infinity and do not necessarily correspond to the (gauge-dependent) coordinate BH positions in the NR simulation. See Ref. [20] for more details.

C. Aligned-spin model, NRSur3dq8Remnant

NRSur7dq4Remnant is restricted to $q \leq 4$ due to a lack of sufficient precessing simulations at higher mass ratios [29]. NR simulations become increasingly expensive as one approaches higher mass ratios and/or spin magnitudes. However, the SXS Catalog has good coverage for aligned-spin BBHs up to $q \leq 8$ [29,31]. We make use of the 104 NR waveforms used for the surrogate model of Ref. [31], which cover the parameter space $q \leq 8$, $|\chi_{1z}|, |\chi_{2z}| \leq 0.8$.

Note that the spins in aligned-spin BBHs are restricted to the \hat{L} direction; this reduces the parameter space to three dimensions. Following Refs. [23,31], we parametrize the NRSur3dq8Remnant fit as follows,

$$\Lambda = [\log(q), \hat{\chi}_a^{\text{coorb}}, \chi_a^{\text{coorb}}], \quad (11)$$

where we use Eqs. (7) and (10) but keep in mind that spins in the coorbital frame are the same as those in the inertial frame for aligned-spin systems.

III. MODELING ERRORS

We evaluate the accuracy of our new surrogate models by comparing against the NR simulations used in this work. To avoid underestimating the errors, we perform a 20-fold cross-validation study to compute “out-of-sample” errors as follows. We first randomly divide the training simulations into 20 groups of roughly the same size. For each group, we build a trial surrogate using the remaining training simulations and test against the simulations in that group, which may include points on the boundary of the training set.

For comparison, we also compute the errors for existing peak luminosity fitting formulas [9,10,17] against the NR simulations. We refer to the fit of Ref. [9] as UIB,³ the fit of Ref. [10] as HL,⁴ and the fit of Ref. [17] as FK.⁵ Note that these fits are not trained on precessing simulations. As the spins evolve for precessing systems, there is an ambiguity about at what time these fits should be evaluated. We follow the procedure outlined in Ref. [40] and used in LIGO/Virgo analyses (e.g., Ref. [7]): NR spins are evolved from relaxation to the Schwarzschild innermost stable circular orbit (ISCO) using PN theory. The spins at ISCO, projected along \hat{L} , are used to evaluate the aligned-spin peak luminosity fitting formulas.

A. Errors for the precessing model

We demonstrate the accuracy of the NRSur7dq4Remnant model by comparing against the 1528 precessing NR simulations described in Sec. II B. We perform a 20-fold cross-validation study where we leave out approximately 75 simulations in each trial for testing. Figure 1 shows the




FIG. 1. Fractional errors (out of sample) in predicting the peak luminosity for the precessing model NRSur7dq4Remnant when compared against precessing NR simulations. When evaluating NRSur7dq4Remnant, we use the NR spins at $t = -100M$, where the model was trained. Also shown are the NR resolution errors and errors for different existing fitting formulas. The square (triangle) markers at the top indicate the median (95th percentile) values. NRSur7dq4Remnant is more accurate than the existing formulas by about an order of magnitude.

errors for NRSur7dq4Remnant when using the NR spins at $t = -100M$ as the input. As the model was trained at this time, these errors represent the errors in the GPR fitting procedure. The 95th percentile fractional error in predicting the peak luminosity is approximately 0.02. We also show the errors for existing fitting formulas and the NR resolution error, estimated by comparing the two highest resolution simulations. Our errors are at the same level as the estimated NR error and about an order of magnitude smaller than that of existing fitting formulas.

In practice, one might want to specify the input spins at arbitrary times. For example, in LIGO-Virgo analyses (e.g., Ref. [3]), the spins are measured at a fixed reference frequency. Following Refs. [20,23], this is handled by evolving the input spins from the reference frequency to $t = -100M$ using a combination of PN in the early inspiral and NRSur7dq4 [20] spin evolution in the late inspiral. Figure 2 shows the errors in NRSur7dq4Remnant when the spins are specified at a reference orbital frequency $f_{\text{ref}} = 20$ Hz. These errors are computed by comparing against 23 long NR ($3 \times 10^4 M$ to $10^5 M$ in length) simulations [20,41] with mass ratios $q \leq 4$ and generically oriented spins with magnitudes $\chi_1, \chi_2 \sim 0.5$. Note that none of these simulations were used to train the surrogates. Comparing with Fig. 1, even with spin evolution, our errors are about an order of magnitude lower than that of existing fits.

B. Errors for the aligned-spin model

We demonstrate the accuracy of the NRSur3dq8Remnant model by comparing against the 104 aligned-spin NR simulations described in Sec. II C. We perform a 20-fold

³After the research group.

⁴For the authors Healy + Lousto.

⁵For the lead authors Forteza + Keitel.




FIG. 2. Fractional errors for NRSur7dq4Remnant in predicting the peak luminosity when spins are specified at a reference frequency of $f_{\text{ref}} = 20$ Hz. For four different total masses, we compute the errors against 23 long NR simulations that were not used to train the model. For each mass, the errors are shown as a smoothed vertical histogram (or a violin). The histograms are normalized so that all violins have equal width.

cross-validation study where we leave out approximately 5 simulations in each trial for testing. These errors are shown in Fig. 3. The 95th percentile fractional error in predicting the peak luminosity is approximately 0.002. Figure 3 also shows the errors for the existing fitting formulas and the estimated NR errors. NRSur3dq8Remnant is comparable to NR and more accurate than existing fits by at least an order of magnitude.

We note that the estimated NR errors for aligned-spin BBHs (Fig. 3) are significantly smaller than that for




FIG. 3. Fractional errors (out of sample) in predicting the peak luminosity for the aligned-spin model NRSur3dq8Remnant when compared against aligned-spin NR simulations. Also shown are the NR resolution errors and errors for different existing fitting formulas. The square (triangle) markers at the top indicate the median (95th percentile) values. NRSur3dq8Remnant is more accurate than the existing formulas by at least an order of magnitude.

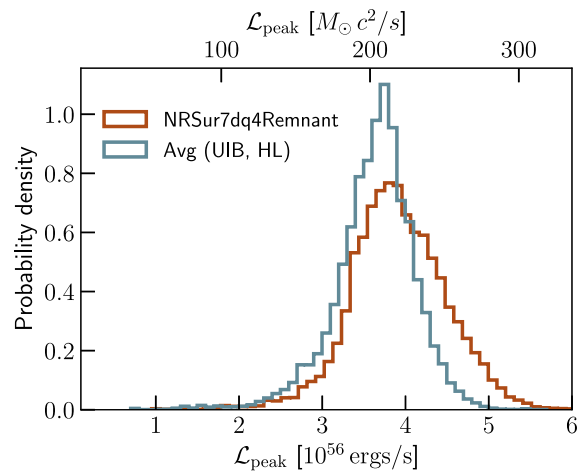


FIG. 4. Posterior distribution for the peak luminosity of GW190521, obtained using the NRSur7dq4Remnant model as well as the average of the UIB [9] and HL [10] fitting formulas. While the two posteriors are consistent with each other, NRSur7dq4Remnant suggests a slightly higher value for the peak luminosity.

precessing BBHs (Fig. 1). The reason for this is not clear, but this places a limit on how accurate the surrogate models can be. This is reflected in the higher errors for NRSur7dq4Remnant compared to NRSur3dq8Remnant. More accurate precessing NR simulations may be necessary to further improve the precessing model.

IV. PEAK LUMINOSITY OF GW190521

As a first application of our models, we compute the peak luminosity of GW190521 [7] using NRSur7dq4Remnant. We apply the NRSur7dq4Remnant model to the posteriors samples for the component masses and spins, obtained using the preferred NRSur7dq4 model in Ref. [8] and made publicly available [42] by the LIGO-Virgo Collaboration. This peak luminosity posterior is shown in Fig. 4. We compare this with the peak luminosity posterior obtained in Ref. [8] using the average of the UIB [9] and HL [10] fitting formulas applied to the same NRSur7dq4 posterior samples. While the two posteriors are consistent with each other, NRSur7dq4Remnant shows support for slightly higher values of peak luminosity. This level of agreement is expected, as GW190521 had a relatively weak signal-to-noise ratio of approximately 14.7 [7]. As GW detectors become more sensitive in the coming years, we can expect to see stronger signals for which systematic biases in peak luminosity models will become important.

V. CONCLUSION

We present GPR based NR surrogate models for peak luminosity of BBH mergers. The first model, NRSur7dq4Remnant, is trained on 1528 precessing systems with mass ratios $q \leq 4$ and spin magnitudes $\chi_1, \chi_2 \leq 0.8$.

The second model, NRSur3dq8Remnant, is trained on 104 aligned-spin systems with mass ratios $q \leq 8$ and spins $|\chi_{1z}|, |\chi_{2z}| \leq 0.8$. Both models are comparable to the NR simulations in accuracy and outperform existing fitting formulae by an order of magnitude or more. The models are made publicly available through the PYTHON module SURFINBH [25] and can be used to estimate the peak luminosity of GW signals. We use NRSur7dq4Remnant to infer the peak luminosity of the GW event GW190521 and find the results to be consistent with previous predictions.

As our GW detectors improve, we will need models that capture the full physics of BBH mergers. NRSur7dq4Remnant is the first peak luminosity model trained on precessing NR simulations. Models such as this will become necessary to accurately infer the peak luminosity as we approach the era of high-precision GW astronomy.

This research made use of data, software and/or web tools obtained from the Gravitational Wave Open Science Center [42], a service of the LIGO Laboratory, the LIGO Scientific Collaboration, and the Virgo Collaboration.

ACKNOWLEDGMENTS

We thank Scott Field, Leo Stein, and Carl-Johan Haster for useful comments. A. T. gratefully acknowledges the support of the United States National Science Foundation for the construction and operation of the LIGO Laboratory and Advanced LIGO, the LIGO Laboratory NSF Research Experience for Undergraduates program, and the Carl A. Rouse Family. V. V. is generously supported by a Klarman Fellowship at Cornell, the Sherman Fairchild Foundation, and NSF Grants No. PHY-170212 and No. PHY-1708213 at Caltech.

-
- [1] J. Aasi *et al.* (LIGO Scientific Collaboration), Advanced LIGO, *Classical Quantum Gravity* **32**, 074001 (2015).
 - [2] F. Acernese *et al.* (Virgo Collaboration), Advanced Virgo: A second-generation interferometric gravitational wave detector, *Classical Quantum Gravity* **32**, 024001 (2015).
 - [3] B. P. Abbott *et al.* (LIGO Scientific, Virgo Collaborations), GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs, *Phys. Rev. X* **9**, 031040 (2019).
 - [4] B. P. Abbott *et al.* (LIGO Scientific, Virgo Collaborations), GW190425: Observation of a compact binary coalescence with total mass $\sim 3.4 M_{\odot}$, *Astrophys. J. Lett.* **892**, L3 (2020).
 - [5] R. Abbott *et al.* (LIGO Scientific, Virgo Collaborations), GW190412: Observation of a binary-black-hole coalescence with asymmetric masses, *Phys. Rev. D* **102**, 043015 (2020).
 - [6] R. Abbott *et al.* (LIGO Scientific, Virgo Collaborations), GW190814: Gravitational waves from the coalescence of a 23 solar mass black hole with a 2.6 solar mass compact object, *Astrophys. J. Lett.* **896**, L44 (2020).
 - [7] R. Abbott *et al.* (LIGO Scientific, Virgo Collaborations), GW190521: A Binary Black Hole Merger with a Total Mass of $150 M_{\odot}$, *Phys. Rev. Lett.* **125**, 101102 (2020).
 - [8] R. Abbott *et al.* (LIGO Scientific, Virgo Collaborations), Properties and astrophysical implications of the $150 M_{\text{sun}}$ binary black hole merger GW190521, *Astrophys. J. Lett.* **900**, L13 (2020).
 - [9] D. Keitel *et al.*, The most powerful astrophysical events: Gravitational-wave peak luminosity of binary black holes as predicted by numerical relativity, *Phys. Rev. D* **96**, 024006 (2017).
 - [10] J. Healy and C. O. Lousto, Remnant of binary black-hole mergers: New simulations and peak luminosity studies, *Phys. Rev. D* **95**, 024037 (2017).
 - [11] J. D. Schnittman, Astrophysics of super-massive black hole mergers, *Classical Quantum Gravity* **30**, 244007 (2013).
 - [12] B. Kocsis and A. Loeb, Brightening of an Accretion Disk Due to Viscous Dissipation of Gravitational Waves During the Coalescence of Supermassive Black Holes, *Phys. Rev. Lett.* **101**, 041101 (2008).
 - [13] G. Li, B. Kocsis, and A. Loeb, Gravitational wave heating of stars and accretion disks, *Mon. Not. R. Astron. Soc.* **425**, 2407 (2012).
 - [14] N. J. Cornish and T. B. Littenberg, BayesWave: Bayesian inference for gravitational wave bursts and instrument glitches, *Classical Quantum Gravity* **32**, 135012 (2015).
 - [15] M. Millhouse, N. J. Cornish, and T. Littenberg, Bayesian reconstruction of gravitational wave bursts using chirplets, *Phys. Rev. D* **97**, 104057 (2018).
 - [16] G. Carullo, G. Riemenschneider, K. W. Tsang, A. Nagar, and W. Del Pozzo, GW150914 peak frequency: A novel consistency test of strong-field General Relativity, *Classical Quantum Gravity* **36**, 105009 (2019).
 - [17] X. Jiménez Forteza, D. Keitel, S. Husa, M. Hannam, S. Khan, L. London, and M. Pürrer, Phenomenological fit of the peak luminosity from non-precessing binary-black-hole coalescences, Tech. Rep. LIGO-T1600018, 2016, <https://dcc.ligo.org/LIGO-T1600018-v4/public>.
 - [18] J. G. Baker, W. D. Boggs, J. Centrella, B. J. Kelly, S. T. McWilliams, and J. R. van Meter, Mergers of non-spinning black-hole binaries: Gravitational radiation characteristics, *Phys. Rev. D* **78**, 044046 (2008).
 - [19] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, Spin-induced orbital precession and its modulation of the gravitational waveforms from merging binaries, *Phys. Rev. D* **49**, 6274 (1994).
 - [20] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, D. Gerosa, L. C. Stein, L. E. Kidder, and H. P. Pfeiffer, Surrogate models for precessing binary black hole simulations

- with unequal masses, *Phys. Rev. Research* **1**, 033015 (2019).
- [21] J. Blackman, S. E. Field, M. A. Scheel, C. R. Galley, C. D. Ott, M. Boyle, L. E. Kidder, H. P. Pfeiffer, and B. Szilágyi, Numerical relativity waveform surrogate model for generically precessing binary black hole mergers, *Phys. Rev. D* **96**, 024058 (2017).
- [22] 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).
- [23] V. Varma, D. Gerosa, L. C. Stein, F. Hébert, and H. Zhang, High-Accuracy Mass, Spin, and Recoil Predictions of Generic Black-Hole Merger Remnants, *Phys. Rev. Lett.* **122**, 011101 (2019).
- [24] V. Varma, M. Isi, and S. Biscoveanu, Extracting the Gravitational Recoil from Black Hole Merger Signals, *Phys. Rev. Lett.* **124**, 101104 (2020).
- [25] V. Varma *et al.*, surfinBH, <https://doi.org/10.5281/zenodo.1418525>.
- [26] M. Boyle and A. H. Mroue, Extrapolating gravitational-wave data from numerical simulations, *Phys. Rev. D* **80**, 124045 (2009).
- [27] The Spectral Einstein Code, <http://www.black-holes.org/SpEC.html>.
- [28] Simulating Extreme Spacetimes, <http://www.black-holes.org/>.
- [29] M. Boyle *et al.*, The SXS Collaboration catalog of binary black hole simulations, *Classical Quantum Gravity* **36**, 195006 (2019).
- [30] SXS Collaboration, The SXS collaboration catalog of gravitational waveforms, <http://www.black-holes.org/waveforms>.
- [31] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, L. E. Kidder, and H. P. Pfeiffer, Surrogate model of hybridized numerical relativity binary black hole waveforms, *Phys. Rev. D* **99**, 064045 (2019).
- [32] K. Barkett, J. Moxon, M. A. Scheel, and B. Szilágyi, Spectral Cauchy-characteristic extraction of the gravitational wave news function, *Phys. Rev. D* **102**, 024004 (2020).
- [33] J. Moxon, M. A. Scheel, and S. A. Teukolsky, Improved Cauchy-characteristic evolution system for high-precision numerical relativity waveforms, *Phys. Rev. D* **102**, 044052 (2020).
- [34] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*, edited by C. E. Rasmussen and C. K. I. Williams (The MIT Press, Cambridge, 2006), ISBN-13 978-0-262-18253-9.
- [35] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, A. Müller, J. Nothman, G. Louppe, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and É. Duchesnay, Scikit-learn: Machine learning in python, *J. Mach. Learn. Research* **12**, 2825 (2012).
- [36] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. Jiménez Forteza, and A. Bohé, Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era, *Phys. Rev. D* **93**, 044007 (2016).
- [37] 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).
- [38] 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).
- [39] E. Poisson and C. M. Will, Gravitational waves from inspiraling compact binaries: Parameter estimation using second postNewtonian wave forms, *Phys. Rev. D* **52**, 848 (1995).
- [40] N. K. Johnson-McDaniel, A. Gupta, P. Ajith, D. Keitel, O. Birnholtz, F. Ohme, and S. Husa, Determining the final spin of a binary black hole system including in-plane spins: Method and checks of accuracy, Tech. Rep. LIGO-T1600168, 2016, <https://dcc.ligo.org/LIGO-T1600168/public>.
- [41] M. Boyle *et al.*, The SXS Collaboration catalog of binary black hole simulations, *Classical Quantum Gravity* **36**, 195006 (2019).
- [42] LIGO Scientific Collaboration and Virgo Collaboration, Gravitational wave open science center, <https://www.gw-open-science.org>.