

# Cost Minimization in Acquisition for Gravitational Wave Surrogate Modeling

Karl Daningburg and Richard O'Shaughnessy  
*Rochester Institute of Technology*

Gravitational wave science is dependent upon expensive numerical simulations, which provide the foundational understanding of binary merger radiation needed to interpret observations of massive binary black holes. The high cost of these simulations limits large-scale campaigns to explore the binary black hole parameter space. Surrogate models have been developed to efficiently interpolate between simulation results, but these models require a sufficiently comprehensive sample to train on. Acquisition functions can be used to identify points in the domain for simulation. We develop a new acquisition function which accounts for the cost of simulating new points. We show that when applied to a 3D domain of binary mass ratio and dimensionless spins, the accumulated cost of simulation is reduced by a factor of about 10.

## I. INTRODUCTION

The Advanced LIGO [1] and Virgo [2] instruments are regularly observing the gravitational waves generated by the merger of two compact objects [3–6]. To determine the parameters of the binaries responsible for these observations, each measurement is compared directly or indirectly with predictions for the gravitational radiation produced when binaries merge [7–9]. Numerical simulations provide the authoritative (and often only) foundational insight into radiation generated during the late stages of mergers [10]. However, these simulations are extremely expensive, making complete exploration of the domain infeasible.

Some scientists have turned to surrogate modeling, a class of mathematical methods which emulate an underlying function. Surrogate models have been widely applied to gravitational waveform modeling [11–22], both to accelerate existing approximations and to interpolate between detailed numerical relativity simulations. Given a gravitational wave model evaluated on candidate binary parameters  $x_k$ , a surrogate model interpolates the waveform's asymptotic amplitude and phase across the binary parameters  $x$  (and time). However, to date these calculations all adopt a simplified way to select the initial training points  $x_k$ : they are usually selected with a greedy algorithm based on a *fast approximation* to the true waveform [16], independent of the cost of simulating those parameters  $x_k$ .

Our novel contribution is to incorporate both interpolator uncertainty and acquisition cost into our acquisition function. This approach balances the goals of (i) accurately modeling the objective function over the entire domain by picking points of maximum variance and (ii) minimizing total acquisition cost by minimizing the cost of selected points. Importantly, our goal is not to create a surrogate model which rivals existing surrogate models in accuracy or speed; it is only to demonstrate the potential of cost-weighted acquisition when it comes to training a surrogate model with Numerical Relativity as cheaply as possible.

This paper is organized as follows. In Section II we define the objective function we have chosen to learn, then

define acquisition functions to be tested. We introduce a workflow for acquiring data and training a Gaussian process. In Section III we examine the performance of the acquisition functions and resulting surrogate models. In Section IV we summarize the work and conclude with future prospects for this research.

## II. METHOD

### A. Surrogate review by example

To be concrete, we introduce surrogate modeling for gravitational waves by summarizing the specific techniques adopted by Blackman et al. in [16] to generate a surrogate gravitational wave model for precessing binaries. The complex dimensionless gravitational wave strain

$$h(t, \theta, \phi; \boldsymbol{\lambda}) = h_+(t, \theta, \phi; \boldsymbol{\lambda}) - i h_\times(t, \theta, \phi; \boldsymbol{\lambda}) \quad (1)$$

is expanded in terms of its polarizations  $h_+$  and  $h_\times$ , where  $t$  is time,  $\theta$  and  $\lambda$  are polar and azimuthal angles of wave propagation from the binary system, and  $\boldsymbol{\lambda}$  is the set of chosen parameters to describe the system. Blackman et al. select black hole spin vectors and initial mass ratio as parameters for their surrogate model; however, the method applies for any parameter of interest. The angular dependence can be efficiently represented with a spin-weighted spherical harmonic decomposition

$$h(t, \theta, \phi; \boldsymbol{\lambda}) = \sum_{l=2}^{\infty} \sum_{m=-l}^l h^{lm}(t; \boldsymbol{\lambda})^{-2} Y_{lm}(\theta, \phi) \quad (2)$$

Following Blackman et al., we will focus only on the  $l = 2, 3$  modes, since higher-order modes have minimal impact on systematic error. Their surrogate models seek to represent each mode function  $h_{lm}(t, \boldsymbol{\lambda})$  versus time and  $\boldsymbol{\lambda}$ .

Blackman et al.'s approach relies on discovering a finite set of greedy parameters

$$G \equiv \{\boldsymbol{\Lambda}_i \in \mathcal{T}\}_{i=1}^N \quad (3)$$

where  $\mathcal{T}$  is a compact region of parameter space. They then perform NR simulations at each greedy parameter, resulting in the greedy solutions  $\{W(t; \boldsymbol{\Lambda}_i)\}_{i=1}^N$ . They transcribe these solutions into a regression model first by constructing an orthonormal linear basis  $B_n = \{e^i(t)\}_{i=1}^N$  spanning the greedy solutions so that

$$W(t; \boldsymbol{\lambda}) \approx \sum_{i=1}^n c_i(\boldsymbol{\lambda}) e^i(t) \quad (4)$$

where coefficient  $c_i$  is the inner product of  $W(t; \boldsymbol{\lambda})$  with  $e^i(t)$ . In principle the greedy parameters  $G$  may be found by iteratively comparing the approximation to the true function  $W(t; \boldsymbol{\lambda})$  and selecting the point of maximum error:

$$E_n(\boldsymbol{\lambda}) = \| W(\cdot; \boldsymbol{\lambda}) - \sum_{i=1}^n c_i(\boldsymbol{\lambda}) e^i(\cdot) \| \quad (5)$$

However, due to the high cost of full numerical simulation, dense evaluation of this error is not feasible. Blackman et al. circumvent this by building a mock surrogate from a cheaper data source: post-Newton waveform models. They assume that the greedy parameters found in construction of the mock surrogate,  $G^{PN}$ , approximate the greedy parameters for the NR surrogate, and thus are able to select training data at a reasonable cost.

Gravitational waves presented as functions of time have complicated dependence on time and binary parameters, which adds additional artificial hurdles when attempting to interpolate them. A powerful technique is to decompose the waveforms obtained by NR simulations into data pieces which vary more slowly and thus are more easily modeled. These disparate models can be combined to form a single surrogate model. Commonly the phase and amplitude of the waves are used as the data pieces; however, for the precessing waveforms which Blackman et al. are targeting, a more sophisticated decomposition is needed. First, each waveform is translated into a coprecessing coordinate frame. Next, a Gaussian filter is used to remove the effect of nutation of the frame's rotational axis. Finally, the waveforms are decomposed into symmetric and antisymmetric amplitudes and phases

$$A_{\pm}^{l,m}(t) = \frac{1}{2} (| \tilde{h}^{l,m}(t) | \pm | \tilde{h}^{l,-m}(t) |) \quad (6)$$

$$\varphi_{\pm}^{l,m}(t) = \frac{1}{2} (\varphi(\tilde{h}^{l,m}(t)) \pm \varphi(\tilde{h}^{l,-m}(t))) \quad (7)$$

where  $\tilde{h}$  is the waveform in the coprecessing frame. The surrogate training data and output are always expressed using this coprecessing-frame amplitude-phase decomposition. Similarly, surrogate models can be constructed for the precessional dynamics themselves, needed to translate between the simulation and coprecessing frame. Fig-

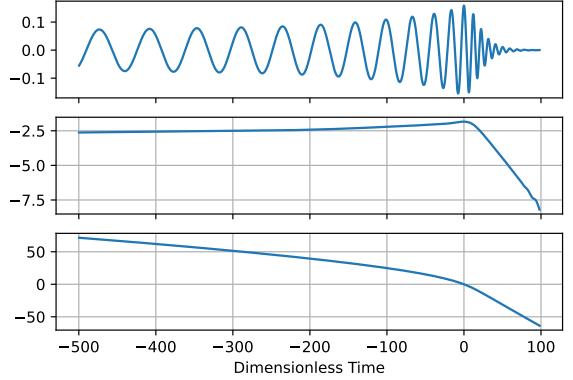


FIG. 1: Strain wave  $h(t)$  is depicted, as well as the wave's decomposition into log amplitude and phase. These smooth functions are more easily interpolated than the raw wave data and thus are used as input to the surrogate model.

ure 1 illustrates the rationale for using an amplitude-phase decomposition for the simple case of a non-precessing binary black hole's dominant mode: in this amplitude-phase decomposition, the two salient functions involved in  $h(t)$  are slowly varying in time (and versus binary parameters).

Now that amplitudes and phases for the decomposed waveforms have been extracted, interpolation on the data can be performed. Blackman et al. use an empirical interpolant via Singular Value Decomposition to unify the  $\tilde{h}(t)$  waveforms into a single surrogate model  $W_S(t; \boldsymbol{\lambda})$ . Despite relying on a SVD algorithm to find basis functions, Blackman et al. use a greedy algorithm to select time points over which to interpolate with these functions. For a detailed discussion on building reduced bases with SVD versus greedy algorithms, see Appendix B of [16]. Blackman et al. use the so-called "empirical time" algorithm developed in [17]. Given a reduced basis  $\{e_i\}_{i=1}^m$  whose span approximates the objective function (in our case, gravitational wave phase and amplitude), the Empirical Interpolation Method (EIM) yields a set of  $m$  empirical times  $T_i$  chosen from the complete set of times  $t_i$ . These times are selected by comparing the value of the empirical interpolant

$$\mathcal{I}_j[h](t; \boldsymbol{\lambda}) = \sum_{i=1}^j C_i(\boldsymbol{\lambda}) e_i(t) \quad (8)$$

to the basis functions evaluated at all times, where the  $C_i$  coefficients are the solutions to the interpolation problem at the empirical times. At each step, one more basis function is added to  $\mathcal{I}_j[h](t; \boldsymbol{\lambda})$ ; the time of maximum error is found and added to the set of empirical times  $T_i$ . This process continues until the number of times chosen is equal to the number of basis functions  $m$ .

Blackman et al. acquire data for their surrogate model by maximizing error between the true function and the

surrogate. Since NR data is far too expensive to obtain in the quantities needed for this method, they evaluate a cheaper objective function to acquire new points. Another type of acquisition function which may be more easily used with an expensive reference function is one based on minimizing uncertainty rather than error. This formulation of the interpolation problem lends itself to Gaussian processes, which provide quantification of uncertainty throughout the domain. Therefore, we propose to focus on Gaussian processes as interpolators and measures of uncertainty as a main ingredient in our acquisition functions.

### B. Gaussian process interpolation

Stochastic processes may be understood as an extension of probability distribution functions from random variables to functions. We utilize the Gaussian process (GP), which by generalizing the Gaussian distribution simplifies computations required for learning [23]. Analogous to the mean and variance parameters of the Gaussian, GPs are fully defined by a mean function,  $\mu(x)$ , and a covariance (or kernel) function,  $k(x, x')$ . GPs can interpolate across a domain, representing the infinite number of possible true functions as variance. Once trained, a GP can be evaluated to provide new data in emulation of the objective function. Gaussian processes have been previously employed to perform the interpolation needed to construct surrogate models and identify candidate points for followup analysis; see, e.g., [22], which is similar to our cost-neutral approach.

The choice of a kernel function is not a straightforward one. Since kernel functions encapsulate assumptions about the function to be learned, prior knowledge of this function is useful. For this introductory work, we use the simple and commonly-used squared exponential (SE) function, with the understanding that further investigation is worthwhile:

$$K(X, Y) = \sigma^2 \exp \left\{ -\frac{(X - Y)^2}{2l^2} \right\} \quad (9)$$

where  $X$  and  $Y$  are the two inputs,  $\sigma^2$  is the variance, and  $l$  is the characteristic length scale, which modifies the correlation between nearby data points. The SE function has the advantage of being easy to evaluate and infinitely differentiable, and it yields smooth interpolation between data points.

Gaussian processes are inherently probabilistic, allowing them to express uncertainty in a given domain as *variance*. The equation for a GP's variance is shown below, where  $K()$  represents the selected kernel function,  $X$  represents the training data, and  $X_*$  is the point at

which variance is computed.

$$V = K(X_*, X_*) - K(X_*, X)[K(X, X)]^{-1}K(X, X_*) \quad (10)$$

### C. Simulation Cost

Although Numerical Relativity (NR) simulations were not used for training this surrogate model, we develop this model with the intent of directly targeting and interpolating NR simulations. As our goal is a reduction of simulation and training cost, we develop a cost function to estimate the expense of simulating points in the domain. Our colleagues estimate [24] the length of a simulation scales approximately with mass ratio  $q$ , as does the resolution required to maintain accuracy; therefore the cost dependence on mass ratio can be modeled as  $q^2$ . Meanwhile, dimensionless spins of the two bodies have very little effect on the simulation cost until spin of 0.6 is reached. Thereafter, the cost increases rapidly and diverges as spin approaches 1. After consulting with an NR group [24], the following piecewise function models the cost-spin relationship:

$$\mathcal{S}(\chi) = \begin{cases} 1 & \chi \leq 0.6 \\ \frac{0.4}{1-\chi} & 0.6 < \chi \leq 0.9 \\ \frac{0.04}{(1-\chi)^2} & 0.9 < \chi < 1 \end{cases} \quad (11a)$$

For our chosen domain of mass ratio and spins for bodies 1 and 2, the cost function is as follows:

$$C(q, \chi_1, \chi_2) = q^2 + \mathcal{S}(\chi_1) + \mathcal{S}(\chi_2) \quad (11b)$$

where  $q$  is mass ratio  $\frac{m_1}{m_2}$ ,  $m_1 > m_2$ , and  $\chi_1$  ( $\chi_2$ ) is dimensionless spin magnitude for body 1 (2).

### D. Acquisition

Acquisition refers to the collection of a data point, whether by experimentation or simulation. In order to train our surrogate model most efficiently, we develop acquisition functions which assign a value to each point in the domain. Optimizers are used to maximize these acquisition functions, and the resulting point in the domain is selected for the next simulation.

Abstractly, we wish to formulate an acquisition function which balances the need to find points of maximum uncertainty with the requirement of keeping simulation costs at a minimum. Having defined variance and cost in Equations (10) and (11b), we can now construct an acquisition function which quantifies our learning per unit

cost:

$$A = \frac{V}{C} \quad (12)$$

Maximizing this function at each step constitutes a greedy algorithm for learning as much about the objective function as possible for the lowest cost.

For the purpose of comparison, we test both a variance maximizing routine (see, e.g., [22]) and a variance-to-cost ratio maximizing routine. For brevity, we will refer to these as  $V$  and  $V : C$  routines, respectively.

### E. Time Selection

The data available for a given point in parameter space is comprised of a complete gravitational waveform which may be expressed in the frequency domain or the time domain. It is desireable to select an optimal set of times (or frequencies) at which to measure the amplitude or phase of each available GW and interpolate that data through parameter space, allowing us to reconstruct any waveform in the domain.

The algorithm for time selection proceeds as follows. Given a set of GWs associated with a set of parameters, we seek a minimal subset of dimensionless time points  $\tau$  to fully characterize the behavior of any waveform in the domain; we will therefore call this set *characteristic time*  $\tau_c$ . To initialize the algorithm, five evenly-spaced time points are selected. The data corresponding to these times are fed to SciPy's `interp1d` routine, which trains a cubic spline interpolator for each available GW. We use this spline to interpolate across 50000 evenly spaced time points from -500 to  $+75\tau$ . The root-mean-square error

$$\text{RMSE}(\phi_j) = \frac{1}{n} \sum_{i=1}^n \sqrt{\left( \frac{\phi_f^{(ij)} - \phi_c^{(ij)}}{\max |\phi_f^{(ij)}|} \right)^2} \quad (13)$$

is computed at each time step across all splines. Here  $\phi$  is the GW phase,  $j$  indicates one of the 50000 time steps,  $n$  is the number of GWs available,  $\phi_f$  is the fiducial spline, and  $\phi_c$  is the spline trained only on  $\tau_c$ . The time of maximum error is chosen to be added to the spline training data. The process iterates until the maximum  $\text{RMSE}(\phi_j)$  falls below a desired threshold. These characteristic time points and the corresponding values of phase or amplitude are used to train Gaussian processes for interpolation across parameter space. New GWs can be constructed by training new splines on the data interpolated by the GPs.

### F. Error Metric

We can evaluate the model's performance at any point in the space  $\mathcal{T}$  by comparing it to the source of our data,

IMRPhenomD. A brief discussion of *mismatch*, a common metric for quantifying the difference between waveforms, is necessary. Given two waveforms, a fiducial wave  $h_1$  and a surrogate approximation  $h_2$ , we require an objective measure of how well  $h_2$  approximates  $h_1$ . For a single mode we compute the *match* [25]:

$$\mathcal{O} = \frac{\max_t \left| \int_{-\infty}^{\infty} h_2^*(f) h_1(f) e^{i2\pi f t} df \right|}{\|h_1\| \|h_2\|} \quad (14)$$

This expression is evaluated using a white-noise detector power spectrum; alternatively, this expression corresponds to the standard equally-weighted Hilbert space inner product in time or frequency. Empirically we estimate the overlap from our discrete data using an inverse Fourier transform

$$\mathcal{O} = \max_t \text{IFT} \left[ N \left( \frac{\tilde{h}_2^*}{\|\tilde{h}_2\|} \right) \left( \frac{\tilde{h}_1}{\|\tilde{h}_1\|} \right) \right] \quad (15)$$

where  $N$  is the number of samples,  $\tilde{h}$  indicates a single wave mode in the frequency domain, and  $*$  represents the complex conjugate. Two waves identical to a constant factor will have an overlap of unity; therefore to reframe as an error metric we compute the mismatch as

$$\mathcal{E} = 1 - \mathcal{O} \quad (16)$$

We chose this approach since this surrogate model is intended to be generic, independent of any detector's power spectrum. For a more complete discussion of waveform mismatch, see [16].

### G. Testing with fast aligned model

Our surrogate modeling methods are similar in principle to those outlined in [16]. Given a physical system parametrized by  $\lambda \in \mathcal{T}$ , where  $\mathcal{T}$  is a compact region in parameter space, we seek cheaply evaluated functions of time  $W(t; \lambda)$  that describe the system. For our model  $\lambda$  includes mass ratio  $q$  and spins ( $\chi_{1z}$  and  $\chi_{2z}$ ) of black holes in a binary black hole system and  $\mathcal{T}$  encompasses a range of mass ratios and spins for which an existing model is available.

Although NR is the eventual data source for this type of model, we begin with a more rapidly solved method: IMRPhenomD [26]. IMRPhenomD is a phenomenological model which evaluates the gravitational wave signal of black hole binaries throughout the inspiral, merger, and ringdown phases. This model is capable of producing data for non-precessing black holes of mass ratio up to 18:1 and dimensionless spin of 0.85, or 0.98 for 1:1 mass systems. IMRPhenomD is a hybrid model which incorporates both Effective-One-Body data and NR simulations for tuning. It has demonstrated error of under 1% versus NR test waveforms [26]. We selected this model to

enable more rapid training of our surrogate model and thorough error analysis of the methods developed. Its close approximation of NR results ensure translation of our methods to NR data acquisition with similar performance.

Once our parameter space is populated following the method outlined in Sec. II D, we turn to the task of interpolating the data across both time and the parameter space. We seek a set of dimensionless times  $\tau_0, \dots, \tau_k$  for each value to be interpolated across parameter space which characterizes the behavior of these values over time for all systems in the target region  $\mathcal{T}$ . Using the method described in Sec. II E, we choose characteristic times  $\tau_c$  for both phase and amplitude of the GWs. Importantly, the times for phase and the times for amplitude are selected separately, since the behavior of these two values is disparate. Once the characteristic times are selected, Gaussian processes are trained on the value of amplitude or phase at each characteristic time for all available gravitational waves.

To solve the surrogate model at any point in parameter space, first we evaluate the Gaussian processes at all characteristic times for both amplitude and phase. Then we reconstruct the full amplitude and phase curves in time using the natural cubic spline interpolant. These curves can be used to reconstruct the original strain waveform  $h(t)$ . To get a succinct understanding of competing models' performance, we evaluate the wave mismatch [see Eqs. (14) and (16)] at select points distributed evenly through parameter space and take the L2 norm to produce a single performance metric. We compute this error metric at strategically chosen intervals of acquisition to understand the overall error trend as more points are acquired.

### III. RESULTS

To compare the performance of two acquisition functions of (a) variance maximization (hereafter  $V$ ) and (b) variance-to-cost ratio maximization (hereafter  $V : C$ ), we devised the following test. For parameter space  $\mathcal{T}$  we chose the three dimensions upon which simulation cost relies: mass ratio  $q$  and the spins  $\chi_1$  and  $\chi_2$ . Mass ratio  $q$  we allowed to vary from 1 to 5, while each  $\chi$  ranged from 0 to 0.99. For the purposes of characteristic time, we chose a maximum L2 norm error of 0.0001 across all acquired waveforms, with a maximum number of characteristic time points of 200. Recall that for each characteristic time point, a Gaussian process will be trained on the data from that moment in time for all acquired waveforms (see subsection II E); therefore the cap of 200 is chosen to keep runtime reasonably low. We initialize the model with 50 data points selected randomly through Latin Hypercube Sampling, and then begin the iterative process of training and acquiring new data.

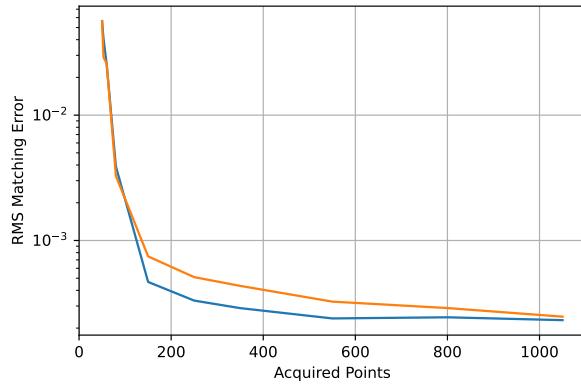


FIG. 2: The mismatch (16) of two surrogate models over data acquisition. In blue the model acquires new points based only on GP variance; in orange, the ratio of variance to cost. The variance-to-cost acquisition function remains competitive with the variance alternative in all but the most high-cost areas.

#### A. Global performance

A comparison of the two surrogate models is presented in Figure 2. For this figure, the mismatch was measured at 125 points evenly spaced in the three dimensional parameter space. Mass ratio  $q$  was sampled in intervals of 1 from 1 to 5, and spins  $\chi_1$  and  $\chi_2$  were sampled in intervals of 0.2 from 0 to 0.8. These samples match the limits of the parameter space with the exception of  $\chi$ , which extend to 0.99. We removed  $\chi = 0.99$  from the global performance calculation due to the relatively high error in this expensive region dominating the L2 error norm of the  $V : C$  routine. By the metric of mismatch vs number of acquired points, the  $V$  scheme slightly outperforms the  $V : C$  scheme until the two converge after 1000 acquisitions.

Reframing the problem as maximizing learning per unit cost rather than per data point acquired reveals the advantage of the  $V : C$  acquisition function. Figure 3 shows the relationship between matching error and relative cost of acquired data points. While performance per cost is similar at higher errors, as the error is driven lower by further data acquisitions the  $V : C$  routine strongly outperforms the  $V$  routine. The  $V : C$  routine achieves the same RMS matching error across the described domain for about an order of magnitude smaller cost.

#### B. Performance over parameter space

The  $V : C$  methods achieve their improved global performance by sacrificing exploration and thus accuracy in the most costly regions. To highlight this tendency, Figure 4 shows the mismatch versus binary parameters after a fixed number of acquisitions. While the  $V : C$ -derived

### C. Discussion

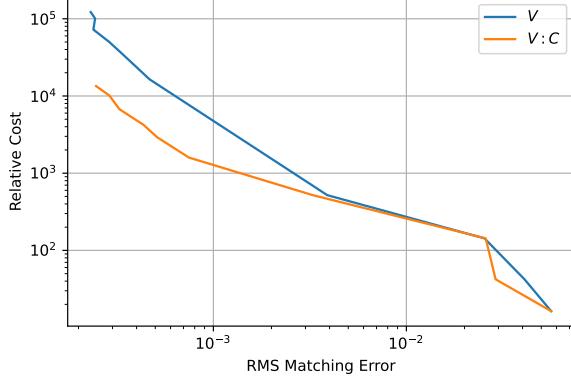


FIG. 3: The relative acquisition cost of two surrogate models over RMS mismatch [Eq. (16)]. In blue the model acquires new points based only on GP variance; in orange, the ratio of variance to cost. The  $V : C$  acquisition function achieves the same RMS match error across the majority of the parameter space with nearly an order of magnitude smaller cost once error is driven below  $10^{-3}$ .

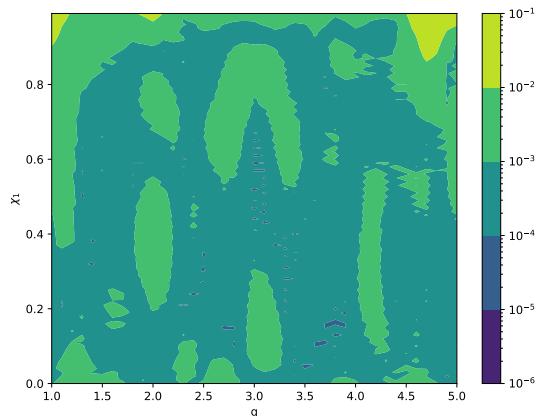


FIG. 4: The mismatch in a slice of parameter space for the  $V : C$  routine after 1000 acquisitions. Note that error remains high in a few ultra-high cost regions. For this slice of parameter space,  $\chi_2$  is set to 0.99.

surrogate model works extremely well in low-cost regions, it does not as reliably reproduce high-cost waveforms.

To date, nature has provided principally low-spin compact binaries, so a focus on low-spin exploration seems warranted. However, high-spin binaries have considerable discovery potential and interesting physics. A more thorough review of the science return of a given GW model based on its accuracy in different regions of parameter space is required to better quantify this tradeoff.

There are several key distinctions between the surrogate model developed here and the model developed by Blackman et al. [16]. Perhaps the most important distinction is the use of a Gaussian process as an interpolant rather than an SVD-based empirical interpolant. Using a Gaussian process opened up the possibility of using the interpolator's uncertainty (variance) about the objective function in order to place simulations to acquire more data efficiently. This approach is simpler and more direct than the method used in [16]; Blackman et al. had to resort to building a mock surrogate to evaluate error between the interpolator and the data source, since NR simulations are too expensive for this type of acquisition. This complicates the model training process, and the use of variance as a stand-in for error circumvents this step and simplifies our algorithm. However, as a non-spectral method, the specific GP interpolation adopted in this work converges relatively slowly versus the number of new simulations acquired. We expect that our technique can be applied to scenarios with other empirically-estimated error metrics like the SVD approach.

The second most important difference between the surrogate described in [16] and the model presented here is incorporation of simulation cost. While both approaches seek to minimize the total number of NR simulations required to train a model, our method gives consideration to the variability in cost of the simulations themselves. While we did not directly compare data acquisition cost between Blackman et al.'s model and our  $V : C$  model, we did use the variance maximization routine  $V$  as a stand-in for this comparison. We found that a cost savings of about tenfold was gained by using the  $V : C$  acquisition function. Of course, due to its extremely rapid (exponential) convergence, the asymptotic behavior of a spectral interpolation should be superior. Our goal in this study is simply to demonstrate the utility of incorporating cost estimates into surrogate simulation targeting.

### IV. CONCLUSION

In this paper, we demonstrate how to incorporate simulation cost into a concrete strategy for iteratively building a surrogate for gravitational waves generated by black hole binaries. By minimizing the estimated (Gaussian process) variance per unit cost, we show how to target synthetic GW simulations to iteratively assemble a comparable-accuracy surrogate at lower cost than a comparable cost-neutral approach. At the mismatch scales appropriate to contemporary gravitational waveform modeling (i.e.,  $10^{-3}$ ), we demonstrated our method should be about an order of magnitude less costly than a comparable cost-neutral approach over the same parameter space.

Though we focused on synthetic gravitational wave simulations specifically, our methods could be directly

transferable to other data-sparse domains with high-cost simulations, including as a wide variety of applications

in astrophysics involving compact object mergers.

---

[1] LIGO Scientific Collaboration, J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, K. Ackley, C. Adams, T. Adams, P. Addesso, et al., *Class. Quant. Grav.* **32**, 074001 (2015), 1411.4547.

[2] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou, A. Allocca, J. Amarni, P. Astone, G. Balestri, G. Ballardin, et al., *Classical and Quantum Gravity* **32**, 024001 (2014), URL <https://doi.org/10.1088/2F0264-9381%2F32%2F2%2F024001>.

[3] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, et al., *Phys. Rev. Lett* **116**, 061102 (2016).

[4] The LIGO Scientific Collaboration, The Virgo Collaboration, B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, et al., *Phys. Rev. X* **9**, 031040 (2019).

[5] The LIGO Scientific Collaboration, the Virgo Collaboration, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. X. Adhikari, et al., *Physical Review X* **11**, 021053 (2021), 2010.14527.

[6] The LIGO Scientific Collaboration, the Virgo Collaboration, B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, V. B. Adya, et al., Available as LIGO-P2000318 (????), URL <https://dcc.ligo.org/LIGO-P2000318>.

[7] J. Veitch, V. Raymond, B. Farr, W. M. Farr, P. Graff, S. Vitale, B. Aylott, K. Blackburn, N. Christensen, M. Coughlin, et al., *Phys. Rev. D* **91**, 042003 (2015), URL <http://link.aps.org/doi/10.1103/PhysRevD.91.042003>.

[8] J. Lange, R. O’Shaughnessy, and M. Rizzo, Submitted to PRD; available at arxiv:1805.10457 (2018).

[9] G. Ashton, M. Hübner, P. D. Lasky, C. Talbot, K. Ackley, S. Biscoveanu, Q. Chu, A. Divakarla, P. J. Easter, B. Goncharov, et al., *ApJS* **241**, 27 (2019), 1811.02042.

[10] J. Centrella, J. G. Baker, B. J. Kelly, and J. R. van Meter, *Reviews of Modern Physics* **82**, 3069 (2010), 1010.5260.

[11] P. Canizares, S. E. Field, J. Gair, V. Raymond, R. Smith, and M. Tiglio, *Phys. Rev. Lett* **114**, 071104 (2015), 1404.6284.

[12] S. E. Field, C. R. Galley, J. S. Hesthaven, J. Kaye, and M. Tiglio, *Physical Review X* **4**, 031006 (2014), 1308.3565.

[13] J. Blackman, S. E. Field, C. R. Galley, B. Szilágyi, M. A. Scheel, M. Tiglio, and D. A. Hemberger, *Phys. Rev. Lett.* **115**, 121102 (2015), URL <http://link.aps.org/doi/10.1103/PhysRevLett.115.121102>.

[14] R. Smith, S. E. Field, K. Blackburn, C. Haster, M. Purrer, V. Raymond, and P. Schmidt, (arXiv:1604.08253) (2016), URL <http://xxx.lanl.gov/abs/arXiv:1604.08253>.

[15] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, D. Gerosa, L. C. Stein, L. E. Kidder, and H. P. Pfeiffer, *Physical Review Research* **1**, 033015 (2019), 1905.09300.

[16] J. Blackman, S. E. Field, M. A. Scheel, C. R. Galley, D. A. Hemberger, P. Schmidt, and R. Smith, *Phys. Rev. D* **95**, 104023 (2017), 1701.00550.

[17] S. E. Field, C. R. Galley, J. S. Hesthaven, J. Kaye, and M. Tiglio, *Phys. Rev. X* **4**, 031006 (2014), URL <https://link.aps.org/doi/10.1103/PhysRevX.4.031006>.

[18] A. Bohé, L. Shao, A. Taracchini, A. Buonanno, S. Babak, I. W. Harry, I. Hinder, S. Ossokine, M. Purrer, V. Raymond, et al., *Phys. Rev. D* **95**, 044028 (2017), 1611.03703.

[19] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Purrer, *Phys. Rev. Lett* **113**, 151101 (2014), 1308.3271.

[20] G. Pratten, C. García-Quirós, M. Colleoni, A. Ramos-Buades, H. Estellés, M. Mateu-Lucena, R. Jaume, M. Haney, D. Keitel, J. E. Thompson, et al., *Phys. Rev. D* **103**, 104056 (2021), 2004.06503.

[21] S. Ossokine, A. Buonanno, S. Marsat, R. Cotesta, S. Babak, T. Dietrich, R. Haas, I. Hinder, H. P. Pfeiffer, M. Purrer, et al., *Phys. Rev. D* **102**, 044055 (2020), 2004.09442.

[22] Z. Doctor, B. Farr, D. E. Holz, and M. Purrer, *Phys. Rev. D* **96**, 123011 (2017), 1706.05408.

[23] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning* (The MIT Press, 2006).

[24] C. Lousto, personal communication (2021).

[25] H. Cho, E. Ochsner, R. O’Shaughnessy, C. Kim, and C. Lee, *Phys. Rev. D* **87**, 02400 (2013), 1209.4494, URL <http://xxx.lanl.gov/abs/arXiv:1209.4494>.

[26] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Purrer, X. J. Forteza, and A. Bohé, *Physical Review D* **93** (2016), ISSN 2470-0029, URL <http://dx.doi.org/10.1103/PhysRevD.93.044007>.

## Acknowledgments

ROS is supported by NSF AST-1909534 and PHY 1912632, 2012057. We thank Hong Qi and Zohey Doctor for helpful feedback.

## Appendix A: Equations

The Jacobian of the variance of a Gaussian Process using a squared exponential kernel was derived for use in a Python optimization routine:

$$J_m(X_*) = \frac{1}{l^2} \left\{ (K^{-1} + (K^{-1})^T) K_* \right\}^T K_* \left( X_*^{(m,:)} - X^{(m,:)} \right) \quad (A1)$$

where  $J_m$  is the  $m$ th element of the Jacobian,  $K$  is the kernel evaluated on the training data,  $K(X, X)$ ,  $K_*$  is the kernel evaluated at a given point,  $K(X_*, X)$ , and the

superscript  $(m, :)$  represents the entire  $m^{th}$  row of a given matrix.