

Improving the scalability of the Gaussian-process error marginalization in gravitational-wave inference

Miaoxin Liu *Department of Physics, National University of Singapore, Singapore 117551, Singapore*

Xiao-Dong Li

*School of Physics and Astronomy, Sun Yat-sen University Zhuhai Campus,
2 Daxue Road, Tangjia, Zhuhai 519082, People's Republic of China
and CSST Science Center for Guangdong-Hong Kong-Macau Great Bay Area,
Zhuhai 519082, People's Republic of China*

Alvin J. K. Chua*

*Department of Physics, National University of Singapore, Singapore 117551, Singapore
and Department of Mathematics, National University of Singapore, Singapore 119076, Singapore*

(Received 18 July 2023; accepted 25 October 2023; published 15 November 2023)

The accuracy of Bayesian inference can be negatively affected by the use of inaccurate forward models. In the case of gravitational-wave inference, accurate but computationally expensive waveform models are sometimes substituted with faster but approximate ones. The model error introduced by this substitution can be mitigated in various ways, one of which is by interpolating and marginalizing over the error using Gaussian-process regression. However, the use of Gaussian-process regression is limited by the curse of dimensionality, which makes it less effective for analyzing higher-dimensional parameter spaces and longer signal durations. In this work, to address this limitation, we focus on gravitational-wave signals from extreme-mass-ratio inspirals as an example, and propose several significant improvements to the base method: an improved prescription for constructing the training set, graphics processing units-accelerated training algorithms, and a new likelihood that better adapts the base method to the presence of detector noise. Our results suggest that the new method is more viable for the analysis of realistic gravitational-wave data.

DOI: [10.1103/PhysRevD.108.103027](https://doi.org/10.1103/PhysRevD.108.103027)

I. INTRODUCTION

The field of gravitational-wave (GW) astronomy has witnessed remarkable progress so far, with the detection of approximately 90 compact binary coalescences (stellar-mass binary mergers) by the LIGO-Virgo-KAGRA Collaboration [1,2]. Future space-based GW detectors operating in the millihertz frequency band, namely LISA [3], TianQin [4,5], and Taiji [6], will lead to the discovery of new kinds of sources such as binary white dwarfs [7,8], massive binary black-hole mergers [9,10], stellar-mass binary inspirals [11,12], and extreme-mass-ratio inspirals (EMRIs) [13–17]. Gravitational waves generated by all of these extreme astronomical events carry unique information, providing novel insights into the physics and astronomy of such phenomena.

To achieve scientific goals in GW astronomy, it is essential to identify and characterize GW signals within

a noisy data stream. The characterization process involves the inference of astrophysical parameters based on a certain GW source model. The accuracy of parameter estimation is constrained by two factors: the statistical error caused by the noise and the theoretical error due to the use of an inaccurate waveform model. It is known that, for both ground-based [18] and space-based GW detectors [19], the statistical error decreases as the signal-to-noise ratio (SNR) increases, while the theoretical error remains constant; this may lead to the exclusion of the true parameter values with high statistical significance.

In GW data analysis, the deliberate incurrence of theoretical error is a common scenario, as it occurs whenever fast approximate models are used in lieu of more accurate but computationally costly models/simulations (e.g., waveforms from numerical-relativity simulations [20,21]). To account for the presence of (known) theoretical error, Gaussian-process regression (GPR) [22,23], a machine-learning technique, has been proposed as a method for interpolating and marginalizing over such an error [24]. The method fits a Gaussian process to

*alvincjk@nus.edu.sg

a small set of precomputed waveform differences between an accurate fiducial model and an approximate one. This process then serves as a prior distribution for the waveform difference, and can be marginalized over in the standard Bayesian likelihood with the approximate model. The GPR marginalized likelihood, which is informed by accurate waveforms, corrects the search under approximate templates and accounts for any residual model inaccuracy with (generally) more conservative error estimates. This method has since been applied in follow-up studies [25–28].

Previous research has demonstrated the potential of the GPR marginalized-likelihood method to mitigate theoretical error in low-dimensional cases [24,25,28]. However, the curse of dimensionality is a major challenge that hinders the use of GPR even in general applications. The number of training points required to cover a parameter space typically increases exponentially with its dimensionality, while the computational complexity of GPR increases cubically with the size of the training set. In the marginalized-likelihood method, this not only slows down the offline training phase but also the online evaluation phase, negating the speed advantage of using approximate templates.

In this study, we propose multiple improvements to the base GPR method that better adapt it to high-dimensional cases; the most notable is the use of Fisher-information-based Latin hypercube sampling (LHS) to generate a more informative training set with fewer points. We illustrate the efficacy of our approach by applying it to EMRI parameter estimation with a representative “accurate” signal model [29] and an artificially constructed “approximate” template model. Even though accurate next-generation EMRI models [30–32] will not be significantly more costly than existing approximate ones (due to recent computational developments [29,33,34]), we choose EMRIs as our example here because the intrinsic information content and computational complexity of their waveforms epitomize most of the difficulties that inhibit the use of the base GPR method.

The remainder of the paper is organized as follows. Section II A provides a brief overview of the marginalized-likelihood method, while Sec. II B introduces the technique of GPR in the context of waveform interpolation. The training of the GPR model using a precomputed set of waveform differences is discussed in Sec. II C, while Sec. II D reviews the construction of the training set as described in previous studies. These sections provide important background information for understanding the application of the GPR method in GW analysis.

In Sec. III, we describe and demonstrate our proposed improvements to the base GPR method. We introduce a parameter rescaling strategy in Sec. III A, and discuss the new hyperparameters to be trained in Sec. III B. We also compare the LHS training-set construction method to the old method in Sec. III C, concluding that the former contains more information than the latter with the same

density. We describe computational enhancements to training in Sec. III D. We present a new form of the GPR marginalized likelihood that properly treats the presence of detector noise in Sec. III E, and discuss an iterative approach to the GPR method in Sec. III F. Our model is tested on data with simulated detector noise in Sec. III G. Finally, we summarize the new techniques implemented in this work and propose possible computational strategies for further extensions in Sec. IV.

II. BACKGROUND: THE BASE GPR METHOD

A. Marginalized likelihood

For a two-channel GW detector, the single source data can be expressed as

$$x(t) = s(t) + n(t), \quad (1)$$

where $s \equiv (s_I, s_{II})$ is the source signal and $n \equiv (n_I, n_{II})$ is the detector noise. In the standard matched-filtering framework, the data are compared against waveform templates $h \equiv (h_I, h_{II})$ that are parametrized by some astrophysical parameters θ , while the detector noise is treated as a Gaussian and stationary stochastic process.

The Bayesian likelihood of the source parameters is thus [35]

$$L \propto \exp\left(-\frac{1}{2}\langle x - h | x - h \rangle\right), \quad (2)$$

where the noise-weighted inner product $\langle \cdot | \cdot \rangle$ on the space of finite-length time series is given by

$$\langle a | b \rangle = 4\text{Re} \sum_{f>0}^{f_N} df \sum_{\chi=I,II} \frac{\tilde{a}_\chi^*(f) \tilde{b}_\chi(f)}{S_{n,\chi}(f)}, \quad (3)$$

with overtilde denoting discrete Fourier transforms, f_N denoting the Nyquist frequency, and $S_{n,\chi}$ denoting the one-sided power spectral density of the channel noise n_χ . The optimal SNR of a waveform template h is given in terms of this inner product as $\sqrt{\langle h | h \rangle}$, while the overlap between two templates is defined as $\langle h_1 | h_2 \rangle / \sqrt{\langle h_1 | h_1 \rangle \langle h_2 | h_2 \rangle}$.

If an accurate template model $h_{\text{acc}}(\theta)$ is used for the analysis, then the source signal s is well described by the model at the actual parameter values θ_{true} , i.e., $s = h(\theta_{\text{true}})$. The maximum-likelihood estimate θ_{ML} does not generally equal θ_{true} due to the presence of detector noise, but the parameter error $\theta_e = \theta_{\text{ML}} - \theta_{\text{true}}$ is purely statistical in nature as it arises only from n , and is thus fully described by the posterior. On the other hand, if an approximate template model $h_{\text{app}}(\theta)$ is used for the analysis, then the parameter error now contains an additional contribution from the difference $h_e = h_{\text{app}} - h_{\text{acc}}$. For high-SNR sources, this theoretical-error term may exceed the

statistical uncertainties described by the posterior, and thus become the limiting factor in obtaining accurate parameter estimates [36].

The bias from theoretical error can be mitigated by specifying a suitable prior probability distribution $p(h_\epsilon)$ for the waveform difference h_ϵ , then marginalizing over h_ϵ in Eq. (2). This ‘‘marginalized likelihood’’ is given by

$$\mathcal{L} \propto \int_W Dh_\epsilon p(h_\epsilon) L_{\text{acc}}, \quad (4)$$

where L_{acc} is Eq. (2) with $h = h_{\text{acc}} = h_{\text{app}} - h_\epsilon$, and W is the space of waveform differences. In [24], Moore and Gair proposed using GPR to define a Gaussian prior distribution, thus allowing the above integral to be analytically approximated (since L_{acc} is also formally Gaussian).

B. Gaussian-process regression

In the GPR method, $h_\epsilon \in W$ may be modeled as a Gaussian process over the parameter space Θ :

$$h_\epsilon(\boldsymbol{\theta}) \sim \mathcal{GP}(\bar{h}_\epsilon, k), \quad (5)$$

where \bar{h}_ϵ is the (vector-valued) mean of the process, and $k(\boldsymbol{\theta}, \boldsymbol{\theta}')$ is the covariance function of the process. Then the set of waveform differences $\{h_\epsilon(\boldsymbol{\theta}_i) \in W | i = 1, 2, \dots, N\}$ corresponding to a small training set of parameter points $\{\boldsymbol{\theta}_i \in \Theta | i = 1, 2, \dots, N\}$ has a Gaussian probability distribution $\mathcal{N}(\bar{h}_\epsilon, \mathbf{K})$ on W^N [28]:

$$p([h_\epsilon(\boldsymbol{\theta}_i)]) = \frac{1}{(2\pi)^N \det \mathbf{K}} \exp\left(-\frac{1}{2} \mathbf{v}^T \mathbf{K}^{-1} \mathbf{v}\right), \quad (6)$$

where the covariance matrix \mathbf{K} and waveform difference vector \mathbf{v} are expressed respectively by

$$[\mathbf{K}]_{ij} = k(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j), \quad (7)$$

$$[\mathbf{v}]_i = h_\epsilon(\boldsymbol{\theta}_i) - \bar{h}_\epsilon. \quad (8)$$

Note that the normalization constant in Eq. (6) is the square of its usual value for a multivariate Gaussian, due to the two independent channels of the process. Also, \mathbf{v} is a deliberate abuse of notation to cast Eq. (6) in the familiar Gaussian functional form; its components are themselves vectors in W equipped with the inner product (3).

The quadratic form in Eq. (6) may be written as

$$\mathbf{v}^T \mathbf{K}^{-1} \mathbf{v} = \text{tr}(\mathbf{K}^{-1} \mathbf{M}), \quad (9)$$

with

$$[\mathbf{M}]_{ij} = [\mathbf{v}\mathbf{v}^T]_{ij} = \langle h_\epsilon(\boldsymbol{\theta}_i) - \bar{h}_\epsilon | h_\epsilon(\boldsymbol{\theta}_j) - \bar{h}_\epsilon \rangle. \quad (10)$$

In [24] and follow-up work, the mean of the process was taken to be the zero vector. Here we use a nonzero but constant mean \bar{h}_ϵ , which is simply chosen to be the mean of the training set of waveform differences $\{h_\epsilon(\boldsymbol{\theta}_i) | i = 1, 2, \dots, N\}$; doing so improves the regression performance at negligible computational cost. We also remove from Eq. (10) the factor of γ that was introduced in Eq. (15) of [28]. This quantity is defined as the (empirical) ratio between the frequency-averaged power spectral densities of the waveform differences and the detector noise, and was added as a ‘‘fudge factor’’ to the base method to prevent the estimate of statistical error from being dominated by the GPR variance when the noise realization is nonzero. Here, we treat the noise in a more principled way by modifying the definition of the marginalized likelihood in the base method; see Sec. III E.

For any new parameter point $\boldsymbol{\theta}$, the enlarged set $\{h_\epsilon(\boldsymbol{\theta}_i), h_\epsilon(\boldsymbol{\theta})\}$ is again normally distributed with mean \bar{h}_ϵ and the covariance matrix

$$\mathbf{K}_* = \begin{bmatrix} \mathbf{K} & \mathbf{k}_* \\ \mathbf{k}_*^T & k_{**} \end{bmatrix}, \quad (11)$$

where

$$[\mathbf{k}_*]_i = k(\boldsymbol{\theta}_i, \boldsymbol{\theta}), \quad k_{**} = k(\boldsymbol{\theta}, \boldsymbol{\theta}). \quad (12)$$

Since $\{h_\epsilon(\boldsymbol{\theta}_i)\}$ is known, the conditional probability distribution of $h_\epsilon(\boldsymbol{\theta})$ given $\{h_\epsilon(\boldsymbol{\theta}_i)\}$ is also Gaussian:

$$p(h_\epsilon(\boldsymbol{\theta})) \propto \frac{1}{\sigma^2} \exp\left(-\frac{1}{2} \frac{\langle h_\epsilon(\boldsymbol{\theta}) - \mu | h_\epsilon(\boldsymbol{\theta}) - \mu \rangle}{\sigma^2}\right), \quad (13)$$

where $\mu(\boldsymbol{\theta})$ and $\sigma^2(\boldsymbol{\theta})$ are given respectively by

$$\mu = \mathbf{k}_*^T \mathbf{K}^{-1} \mathbf{v} + \bar{h}_\epsilon, \quad (14)$$

$$\sigma^2 = k_{**} - \mathbf{k}_*^T \mathbf{K}^{-1} \mathbf{k}_*. \quad (15)$$

Note that $\mathbf{K}^{-1} \mathbf{v}$ in Eq. (14) and \mathbf{K}^{-1} in Eq. (15) have nothing to do with $\boldsymbol{\theta}$, and can thus be precomputed. The GPR mean $\mu(\boldsymbol{\theta})$ is essentially an interpolant for $h_\epsilon(\boldsymbol{\theta})$, with associated (squared) error given by the GPR variance $\sigma^2(\boldsymbol{\theta})$. We may thus define a new GPR-informed template model as

$$h_{\text{GPR}} = h_{\text{app}} - \mu, \quad (16)$$

which approximates h_{acc} since $h_{\text{acc}} = h_{\text{app}} - h_\epsilon$. Equation (13) also provides the prior for h_ϵ in Eq. (4), which evaluates to the GPR marginalized likelihood (of the base method):

$$\mathcal{L} \propto \frac{1}{1 + \sigma^2} \exp\left(-\frac{1}{2} \frac{\langle x - h_{\text{GPR}} | x - h_{\text{GPR}} \rangle}{1 + \sigma^2}\right). \quad (17)$$

C. Training the Gaussian process

The waveform difference model (5) is specified by the (fixed) mean of the training set and the covariance function k ; the latter depends on hyperparameters that are determined by fitting the Gaussian process to the training set. Previous studies in the GW field have demonstrated that the GPR interpolant and the marginalized-likelihood function exhibit consistent performance across various common choices for k [25]. In this work, we use the squared-exponential covariance function

$$k(\boldsymbol{\theta}, \boldsymbol{\theta}') = \sigma_f^2 \exp\left(-\frac{1}{2}\boldsymbol{\tau}^2\right), \quad (18)$$

with

$$\boldsymbol{\tau}^2 = g_{ab}[\boldsymbol{\theta} - \boldsymbol{\theta}']^a [\boldsymbol{\theta} - \boldsymbol{\theta}']^b, \quad (19)$$

where the hyperparameters consist only of an overall scale factor σ_f^2 and the (independent) components g_{ab} of some constant parameter-space metric \mathbf{g} on $\boldsymbol{\theta}$.

As the size of the training set grows, the covariance matrix \mathbf{K} tends to become ill-conditioned. However, it is common practice to add noise to the training set, which allows for some error in the GPR fit. We transform

$$[\mathbf{K}]_{ij} \rightarrow [\mathbf{K}]_{ij} + \sigma_f^2 \sigma_n^2 \delta_{ij}, \quad (20)$$

where δ_{ij} is the Kronecker delta, and the fractional noise variance σ_n^2 of training-set points is taken to be uniform and fixed. The introduction of noise has the side effect of reducing the condition number of \mathbf{K} for more robust numerical calculations. In this work, we use an empirically determined value of $\sigma_n^2 = 10^{-2}$ throughout.

The Gaussian process is fit to the training set by maximizing (the logarithm of) Eq. (6) as a function of the hyperparameters, i.e., the ‘‘hyperlikelihood’’ Z :

$$\ln Z = -\frac{1}{2}\text{tr}(\mathbf{K}^{-1}\mathbf{M}) - \ln \det \mathbf{K} + \text{const.} \quad (21)$$

Part of this maximization may be done analytically, as $\ln Z$ with $\mathbf{K} = \sigma_f^2 \hat{\mathbf{K}}$ is maximized over σ_f^2 at

$$\sigma_f^2 = \frac{1}{2N} \text{tr}(\hat{\mathbf{K}}^{-1}\mathbf{M}). \quad (22)$$

Substituting Eq. (22) back into Eq. (21), we may instead maximize the scale-invariant log-hyperlikelihood

$$\ln Z = -N \ln \text{tr}(\mathbf{K}^{-1}\mathbf{M}) - \ln \det \mathbf{K} + \text{const.} \quad (23)$$

over the metric components only, which reduces the dimensionality of the hyperparameter space by one.

D. Fisher-coordinate training grid

The waveform derivative ∂h and Fisher information matrix Γ are defined respectively as

$$[\partial h]_a = \frac{\partial h}{\partial [\boldsymbol{\theta}]^a}, \quad [\Gamma]_{ab} = \langle [\partial h]_a | [\partial h]_b \rangle. \quad (24)$$

Let $\{(\lambda_i, \hat{\mathbf{v}}_i)\}$ denote the eigensystem of Γ for the SNR-normalized waveform difference, $h_\epsilon / \sqrt{\langle h_\epsilon | h_\epsilon \rangle}$, evaluated at some reference parameter point. One can then define a new coordinate system centered on that reference point, by taking the semiprincipal axes $\{\lambda_i^{-1/2} \hat{\mathbf{v}}_i\}$ of the associated covariance hyperellipse as basis vectors. A local grid-based training set may be constructed by uniformly placing points on a grid defined by the basis vectors in these ‘‘Fisher coordinates.’’ Previous research [28] has employed this grid-based design.

The main challenge of using a grid-based training set is the issue of scalability in high-dimensional parameter spaces. In this study, we adopt an alternative sampling method and compare it to the traditional grid-based training set in Sec. III C. Subsequently, a training set utilizing the alternative sampling method (within a hyperellipse) is employed as our final model in Sec. III G.

III. IMPROVEMENTS AND RESULTS

As mentioned in Sec. I, we choose the example of an EMRI signal to showcase our improvements to the base GPR method, and to demonstrate the scalability of our results to the typical length and complexity of EMRI waveforms. Throughout this study, the fiducial model is taken as the augmented analytic kludge [26,37] with 5PN-adiabatic evolution [38,39] (5PN AAK), which is publicly available as part of the Fast EMRI Waveforms software package [29]. This choice is motivated not only by the improved realism of the model relative to previous kludges, but also by its computational efficiency (which provides a tractable fiducial likelihood for comparison with the marginalized likelihood).

To construct an approximate model, we artificially modify the time evolution of the slowly evolving orbital parameters (p, e, Y) (the quasi-Keplerian semilatus rectum, eccentricity, and cosine of the inclination) by linearly interpolating between 5PN- and 4PN-adiabatic evolution [40]:

$$\dot{p} = (1 - c)\dot{p}_{5\text{PN}} + c\dot{p}_{4\text{PN}}, \quad (25)$$

$$\dot{e} = (1 - c)\dot{e}_{5\text{PN}} + c\dot{e}_{4\text{PN}}, \quad (26)$$

$$\dot{Y} = (1 - c)\dot{Y}_{5\text{PN}} + c\dot{Y}_{4\text{PN}}, \quad (27)$$

where c is a tunable quantity that is fixed to 0.0001 in this study. (When $c = 0$, the orbital evolution reduces to 5PN; when $c = 1$, it is equivalent to 4PN.) This construction

allows for an approximate model that retains a physically motivated dependence on the EMRI parameters, while producing waveforms that have a controllable overlap with those from the fiducial model.

In this work, an EMRI with redshifted component masses $(\mu, M)_{\text{true}} = (10^1, 10^6)M_{\odot}$, dimensionless spin parameter $s_{\text{true}} = 0.9$, and initial orbital parameters $(p_0, e_0, Y_0)_{\text{true}} = (6.97, 0.1, 0.54)$ is considered as a generic example source. Other source parameters are chosen such that the fiducial signal and the approximate signal have an overlap of 0.84. For simplicity, the long-wavelength approximation for the LISA response [41], $h \equiv (h_{\text{I}}, h_{\text{II}})$, is used instead of full time-delay interferometry [42–44]. The signal is six months long and sampled at 0.2 Hz, while the source distance is adjusted to yield a high but feasible SNR of 100. The GPR marginalized likelihood Eq. (17) is used to estimate six source parameters, $(\mu, M, a, p_0, e_0, Y_0)_{\text{true}}$, assuming all other parameters are known and fixed at their true values. The first application of the base GPR method to EMRIs [28] considered only up to a two-month-long signal and three estimated parameters.

To decrease the computational expense involved in initializing and evaluating the marginalized likelihood, we employ a band-pass filter as done in [28]. This filter is applied to restrict both the data and the templates to the frequency range 3.3–8.3 mHz, outside of which little signal information is present.

A. Rescaled Fisher coordinates

The construction of the training set is based on the Fisher matrix Γ for the SNR-normalized waveform difference rather than the Fisher matrix for the accurate waveform, since the former more closely approximates the optimal hyperparameter metric (which is SNR independent); see Sec. III in [28] for a more detailed discussion. Thus for SNR values > 1 , the bulk of the likelihood is typically comfortably contained within the span of the training set. However, the covariance hyperellipses associated with both matrices can occasionally still be comparable in scale, especially in the minor directions (corresponding to the largest Fisher eigenvalues). When sampling from the likelihood, the coverage of the training set might thus be insufficient in these directions.

To address this, we adopt a strategy of rescaling the basis vectors of the Fisher coordinates as $\{\lambda_i^{-1/2}\hat{\mathbf{v}}_i\} \rightarrow \{f_i\lambda_i^{-1/2}\hat{\mathbf{v}}_i\}$, so as to boost the training-set coverage in the minor directions. An appropriate choice of f_i would depend on the discrepancy between the covariance hyperellipses for the waveform difference and the accurate waveform (essentially, the former should be rescaled such that it contains the latter). As this discrepancy is model and signal specific, such a procedure is necessarily somewhat *ad hoc*. In this study, we perform the rescaling by hand, with the following empirically determined values:

$$\begin{aligned} f_1 &= 8.0, \\ f_2 &= 4.0, \\ f_3 &= 2.0, \\ f_4 &= 1.0, \\ f_5 &= 0.5, \\ f_6 &= 0.5, \end{aligned} \quad (28)$$

where the index is sorted in order of decreasing eigenvalue magnitude.

More explicitly, for a training set centered at $(c_1, c_2, c_3, c_4, c_5, c_6)$ in the parameter coordinates, the point $(x_1, x_2, x_3, x_4, x_5, x_6)$ in the rescaled Fisher coordinates corresponds to

$$\begin{bmatrix} c_1 \\ \vdots \\ c_6 \end{bmatrix} + [f_1\lambda_1^{-1/2}\hat{\mathbf{v}}_1 \quad \dots \quad f_6\lambda_6^{-1/2}\hat{\mathbf{v}}_6] \begin{bmatrix} x_1 \\ \vdots \\ x_6 \end{bmatrix} \quad (29)$$

in the parameter coordinates. To sum up, our rescaling strategy ensures that the resulting training set sufficiently covers the parameter region of interest, so as to avoid inaccurate inference due to regression error.

B. Fisher-coordinate metric hyperparameters

The number of training hyperparameters required to specify the covariance metric in Eq. (19) scales approximately with the square of the parameter-space dimensionality, which again is a challenge to the scalability of GPR. A common approach to mitigating this in many GPR applications is to use a diagonal metric. Here, we instead use the Fisher matrix Γ to place constraints on the metric g_{ab} in Eq. (19), since the former is a good approximation to the optimal values for the latter. Specifically, given the unit eigenvectors $\hat{\mathbf{v}}_i$ and eigenvalues λ_i of Γ , we demand that

$$g_{ab} = [\hat{\mathbf{v}}_1 \quad \dots \quad \hat{\mathbf{v}}_6] \begin{bmatrix} \frac{1}{w_1^2}\lambda_1 & & & \\ & \ddots & & \\ & & \ddots & \\ & & & \frac{1}{w_6^2}\lambda_6 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{v}}_1 \\ \vdots \\ \hat{\mathbf{v}}_6 \end{bmatrix}^T, \quad (30)$$

where now only the w_i are trained hyperparameters. Since σ_f^2 is fixed by Eq. (22), the number of hyperparameters is simply the dimensionality of the parameter space. Restricting the covariance function in this way has a negligible impact on regression performance (relative to training the full metric), while significantly reducing the computational cost of training.

C. Latin hypercube sampling

Although a grid-based construction of the training set is simple to implement, it has a couple of important drawbacks. The first one is that the entropy of a grid-based training set is generally lower than that of a more irregularly distributed training set with the same number of points; i.e., it contains less information [45,46]. The second drawback of using a grid-based training set in higher dimensions is that a larger number of points will lie in low-likelihood regions, since typical likelihoods have a radial falloff in density from the maximum-likelihood point.

To address the first drawback, we adopt LHS [47] as an alternative method of choosing training points. This technique allows a random placement of sample points within a hypercube such that no two samples are aligned (up to a regular partition of the hypercube) in any coordinate direction. In two dimensions, this is equivalent to the classic problem of placing nonattacking rooks on a chess board. We adopt a maximum-distance design for generating the Latin hypercube samples, as proposed by [48]. This approach aims to maximize the distance between all pairs of samples, while minimizing the number of pairs that are separated by the same distance [49]. Thus it prevents highly clustered sample regions, and ensures a more homogeneous distribution of the samples. This variant of LHS is implemented using the Surrogate Modeling Toolbox [50].

In order to compare the performance of an LHS training set against a grid-based training set in GPR, we construct a Fisher-based grid of $N = 4^6 = 4096$ points centered around the signal parameters θ_{true} (such that the grid spans a six-dimensional hypercube of side-length 3 in the rescaled Fisher coordinates), as well as an N -sample LHS training set within the same region. The true source parameters are not included in either set of points. From the Gaussian assumption (6), the entropy of each training set is given by

$$\begin{aligned} H([h_\epsilon(\theta_i)]) &= - \int Dh_\epsilon P([h_\epsilon(\theta_i)]) \ln P([h_\epsilon(\theta_i)]) \\ &= -\mathbb{E}[\ln \mathcal{N}(\bar{h}_\epsilon, \mathbf{K})] \\ &= N(1 + \ln(2\pi)) + \ln \det \mathbf{K}. \end{aligned} \quad (31)$$

Under the same initial hyperparameter values $w_1 = \dots = w_6 = 1$, the entropy of the grid based is smaller than that of the LHS training sets by 1651. The smaller entropy value for the grid-based training set indicates that it contains less information, which turns out to be insufficient for effective training.

Although the hyperlikelihood for the grid-based training set increases asymptotically toward some optimal value during training, the Gaussian process fails to fit the waveform difference adequately, with the GPR error (15) at most




FIG. 1. The training evolution of (the logarithm of) the normalized GPR variance σ^2/σ_f^2 , evaluated at the center of the training region (the true parameters for the source).

evaluation points of interest taking on its maximal value [σ_f^2 in Eq. (22)]. This is illustrated by the top plot in Fig. 1, which shows how the GPR variance (normalized by σ_f^2) at the true signal parameters fails to improve as training proceeds. In contrast, the same plot for the LHS training set tends toward a minimal value $\ll 1$, indicating that the set is more optimal for regression while having the same span and number of points as the grid-based training set.

As for the second drawback of using a grid-based training set, the larger relative volume contained in the “corners” of the hypercube leads to a larger proportion of uninformative points in the set, which adds unnecessary computational cost to both the training and evaluation of the GPR model. To make this intuitive, consider a hypersphere (representing the bulk of the likelihood density) that is inscribed in a hypercube (representing the span of the training grid), in d dimensions. When $d = 2$, the volume outside the hypersphere is 21% of the hypercube volume; this rises to $\geq 92\%$ when $d \geq 6$.

To address this drawback, we implement a further hyperspherical truncation of the LHS training set in our final model (used in Sec. III G). In the rescaled Fisher coordinates, the covariance hyperellipse associated with the Fisher matrix is a hypersphere; we simply enlarge this such that it is inscribed in the hypercube spanned by the grid, and then remove all LHS points lying outside the enlarged hypersphere. In this way, the number of model evaluations in low-likelihood regions is greatly reduced. Both the grid-based training set and the truncated LHS training set are compared visually in Fig. 2.

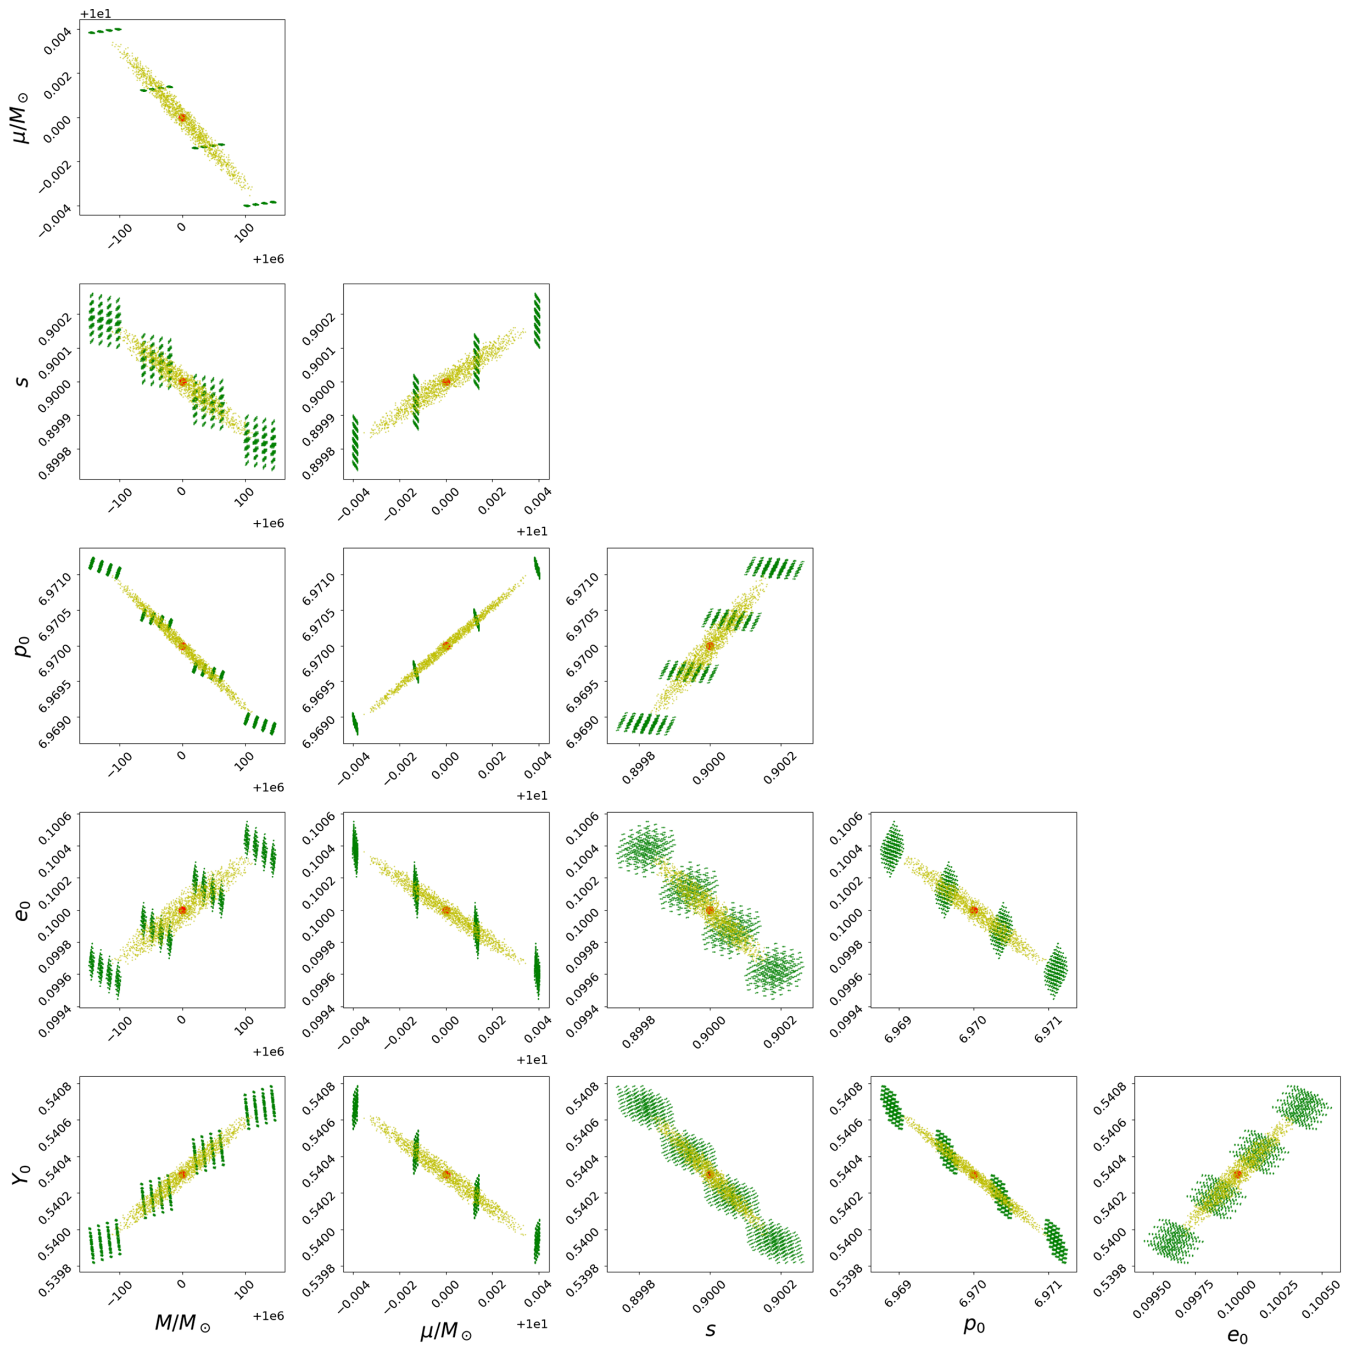


FIG. 2. Visualization of the grid-based training set in Sec. III C (4096 green points) and the hyperspherical truncation of the LHS training set with the same span (1611 yellow points), both centered on the true parameters for the source (red point).

D. Computational acceleration of training

When training the Gaussian process, the cost of evaluating the hyperlikelihood is dominated by the calculation of $\mathbf{K}^{-1}\mathbf{M}$ at each iteration, especially for large N . The most efficient way of computing this quantity is then (i) to solve the linear systems of equations $\mathbf{K}\mathbf{X} = \mathbf{M}$ for \mathbf{X} (instead of inverting \mathbf{K}), and (ii) to parallelize the calculation by performing it on a graphics processing units (GPU). In previous work [25,28], the gains from this approach were

marginal even for the largest considered training sets with $N \sim 10^2$. In this work, where $N \gtrsim 10^3$, it becomes essential. We use the conjugate gradient method [51] to solve for the roots of $\mathbf{K}\mathbf{X} - \mathbf{M}$; this is an iterative technique that is better suited to large N than previously employed methods such as Cholesky decomposition. For a training set containing 4096 points, a single training iteration typically takes around 3.5 seconds when evaluated on a GPU, and around 200 iterations in total to converge.

GPUs can also be used to accelerate evaluation of the trained model [now with fixed \mathbf{K} , such that $\mathbf{K}^{-1}\mathbf{v}$ in Eq. (14) and \mathbf{K}^{-1} in Eq. (15) can be precomputed]. A single evaluation of the GPR mean and variance takes around 0.01 s on a GPU, as compared to around 0.1 s on a CPU. However, note that the marginalized likelihood requires evaluation of the approximate waveform, which must also be accelerated in order to gain the full benefit from the Gaussian-process component of the model. In this work, as the cost of waveform generation is the computational bottleneck in the case of EMRIs, we do not implement the sampling of likelihoods on a GPU.

E. Marginalized likelihood for nonzero noise

From standard noise properties, the expectation of the logarithm of Eq. (17) is given (up to a constant) by

$$E[\ln \mathcal{L}] = -\frac{1}{2} \frac{\langle s - h_{\text{GPR}} | s - h_{\text{GPR}} \rangle + \mathcal{N}}{1 + \sigma^2} + \ln \left(\frac{1}{1 + \sigma^2} \right), \quad (32)$$

where $\mathcal{N} := E[\langle n | n \rangle]$ is the expected noise power. It is the interplay between $\mathcal{N} \neq 0$ and the size/variation of σ^2 over the signal space that is generally problematic for the practical application of the GPR marginalized likelihood with nonzero noise. Essentially, the profile of the likelihood becomes driven by the variation of σ^2 if \mathcal{N} is too large, and it can even have a narrowed credible region that excludes the true parameters to high significance.

This issue was recognized and addressed in [28], although not explicitly described in that paper. There, σ^2 was reduced by an overall factor of $\gamma \ll 1$, with the value of γ chosen empirically as the ratio between the typical power of the waveform differences and the expected power of the detector noise. This works simply because the GPR likelihood approaches the accurate likelihood as $\gamma \rightarrow 0$, but it is rather *ad hoc* and does not generally yield broadened credible regions in the former.

We will instead redesign the GPR likelihood in a way that aims to recover, for general noise realizations, its behavior when $\mathcal{N} = 0$ (while reducing to the accurate likelihood as $\sigma^2 \rightarrow 0$). It is straightforward to achieve this when the likelihood is approximated by Eq. (32); one such solution is simply

$$\mathcal{L} \propto \frac{1}{1 + \sigma^2} \exp \left(-\frac{1}{2} \frac{\langle x - h_{\text{GPR}} | x - h_{\text{GPR}} \rangle + \mathcal{N} \sigma^2}{1 + \sigma^2} \right). \quad (33)$$

Here, \mathcal{N} needs to be estimated accurately because the noise-corrected likelihood can still be quite sensitive to any residual noise power near the maximum-likelihood point. We propose an iterative method for estimating \mathcal{N} (and for refining inference) in the following section.

F. Iterative inference

When applying the GPR method to realistic inference, the starting point is an estimate of the true source parameters θ_{true} , so as to construct the training set in its local vicinity. This estimate is most naturally obtained through maximum likelihood or maximum *a posteriori* estimation with the approximate waveform model (since the accurate model is assumed to be computationally intractable); we denote it by θ_{app} . At this stage, a first estimate of \mathcal{N} is given by

$$\mathcal{N}_{1\text{stGPR}} = \langle x - h_{\text{app}}(\theta_{\text{app}}) | x - h_{\text{app}}(\theta_{\text{app}}) \rangle. \quad (34)$$

Together with the training set centered on θ_{app} and the Gaussian process that is trained on this set, inference can then be performed using the noise-corrected GPR likelihood (33), which we denote by $\mathcal{L}_{1\text{stGPR}}$.

Computation of the posterior under the first GPR model yields a first maximum *a posteriori* estimate, which we denote by $\theta_{1\text{stGPR}}$ (σ^2 taken as 0 in this optimization). Depending on the distance between $\theta_{1\text{stGPR}}$ and θ_{true} (equivalently, the error in the approximate model), the first GPR posterior may not be a sufficiently faithful approximation to the accurate posterior. In this case, another iteration of inference may be performed by constructing a second training set centered on $\theta_{1\text{stGPR}}$, retraining the Gaussian process, and reestimating the noise as

$$\mathcal{N}_{2\text{ndGPR}} = \langle x - h_{\text{acc}}(\theta_{1\text{stGPR}}) | x - h_{\text{acc}}(\theta_{1\text{stGPR}}) \rangle. \quad (35)$$

Here we have used the accurate waveform model to compute the second noise estimate, although we could also have used $h_{1\text{stGPR}}$ instead. In realistic scenarios, such an iterative usage of the GPR method for high-precision inference is unlikely to take more than two iterations; if the error in the approximate model is so large as to require this, a more prudent approach in the first place would be to improve the approximate model or construct better fits to the accurate model.

G. Results

In this subsection, we present illustrative results for the GPR method with all of the proposed improvements in Secs. III A–III E, when applied to our example EMRI signal with simulated LISA noise [52]. We assume a flat prior in a suitably bounded region of parameter space, and employ the Markov chain Monte Carlo sampler emcee [53,54] to draw samples from the standard approximate likelihood L_{app} , the standard accurate likelihood L_{acc} , and the GPR marginalized likelihoods \mathcal{L} . Note that the accurate likelihood is assumed to be unavailable in the actual scenarios to which the GPR likelihood might be applied, but we include it here to showcase the efficacy of the method.

Figure 3 displays the GPR likelihood that is computed in the first iteration described in Sec. III F, along with the standard accurate and approximate likelihoods for comparison. The approximate likelihood excludes the true source parameters at more than 3-sigma significance.

On the other hand, the first GPR likelihood provides a decent (but still slightly shifted) approximation to the accurate likelihood, and agrees with the true parameters to within 3-sigma. Finally, results from a second iteration are presented in Fig. 4, where the second GPR likelihood




FIG. 3. Projected one- and two-dimensional plots comparing the accurate and approximate likelihoods to the first noise-corrected marginalized likelihood (with a truncated LHS training set that is centered at the maximum-likelihood estimate of the approximate waveform). The shown contour level is 3-sigma. The true source parameters are indicated by red dots.

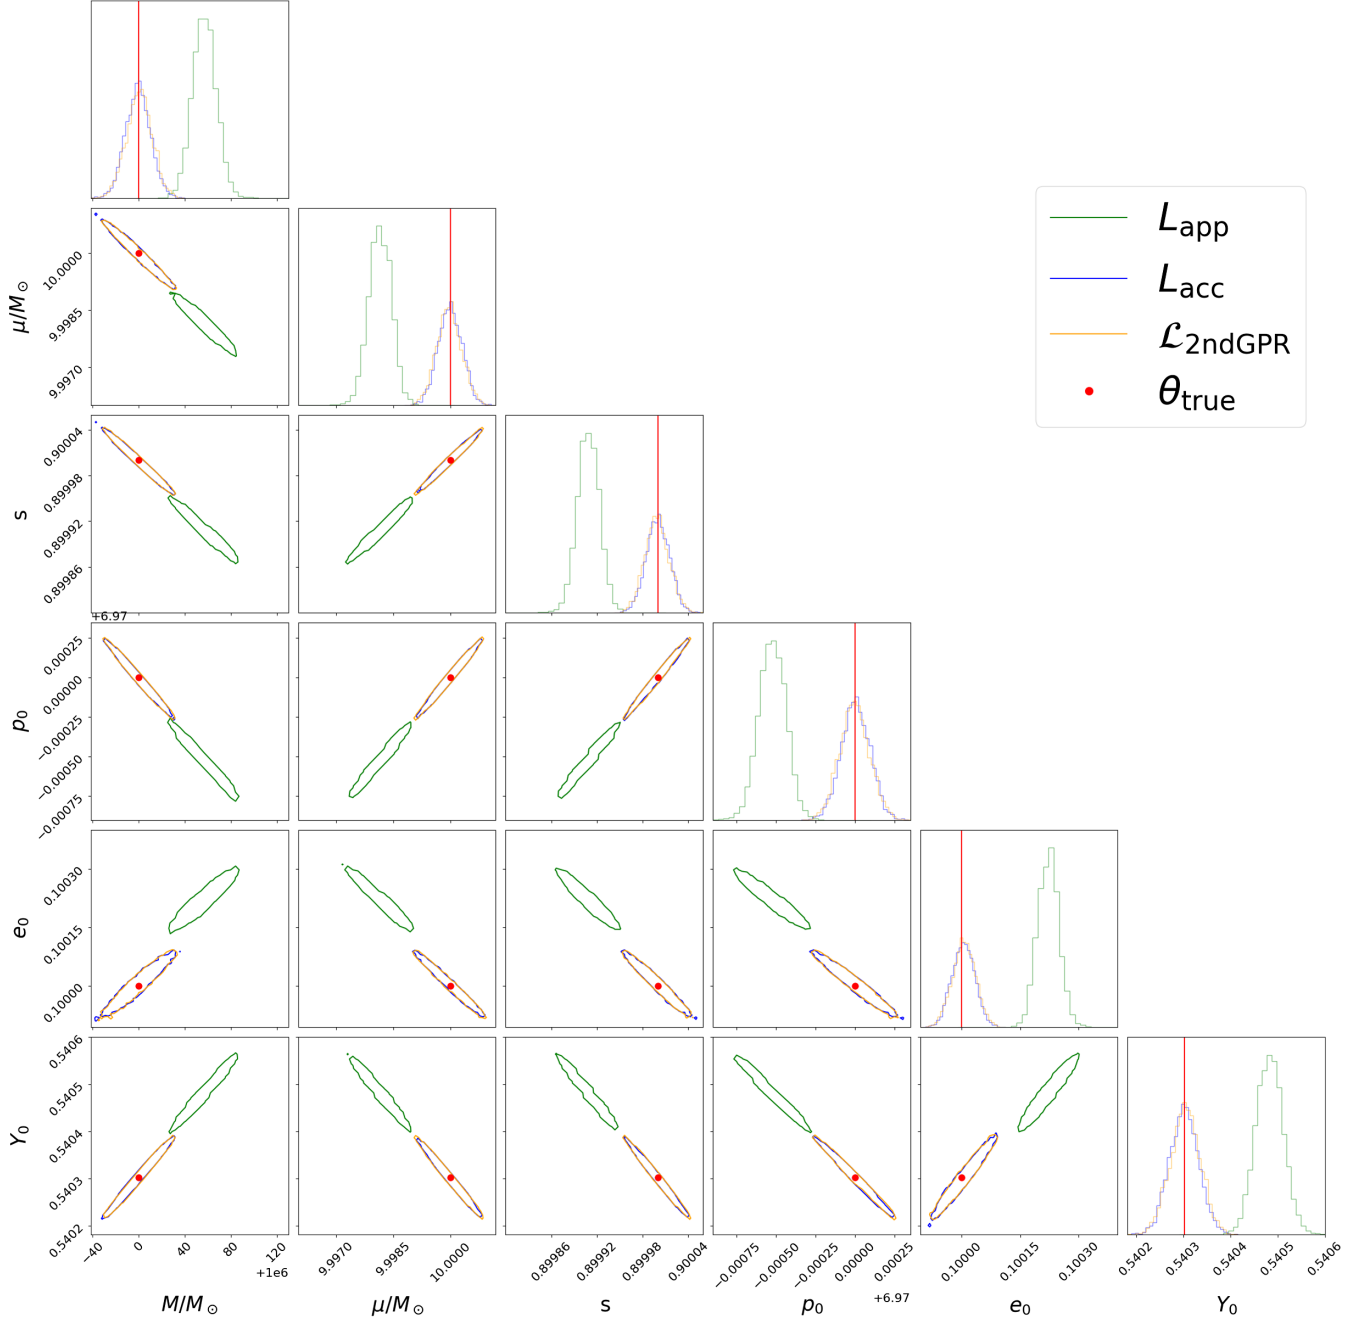


FIG. 4. Projected one- and two-dimensional plots comparing the accurate and approximate likelihoods to the second noise-corrected marginalized likelihood (with a truncated LHS training set that is centered at the maximum-likelihood estimate of the first GPR model). The shown contour level is 3-sigma. The true source parameters are indicated by red dots.

is seen to be almost perfectly consistent with the accurate likelihood.

IV. CONCLUSION

This work improves the scalability of the GPR marginalized-likelihood scheme for high-precision GW inference [24,25,28], thus extending its potential application to higher-dimensional parameter spaces and longer

signal durations (six intrinsic parameters and six-month-long signals, in our EMRI example). Several significant modifications have been made to the base GPR method that was developed in previous work.

In Secs. III A–III D, various improvements to the training of the Gaussian process are described. These are (i) a rescaling of the Fisher-informed training set that is better adapted to highly correlated parameters, (ii) a Fisher-informed constraint on the covariance metric such that

the number of hyperparameters scales linearly rather than quadratically with the parameter-space dimensionality, (iii) the use of LHS and (Fisher-informed) hyperspherical truncation to construct the training set, and (iv) computationally efficient training through the use of the conjugate gradient method and GPU acceleration. These modifications boost the scalability of the GPR method by significantly reducing the required density of the training set in a given region of interest (such that it grows subexponentially with the parameter-space dimensionality), and by greatly accelerating the training process as well.

Sections III E and III F describe improvements to the marginalized-likelihood method itself. We make a crucial redefinition of the marginalized likelihood in order to render it usable in realistic inference scenarios with non-zero detector noise; this is done via an estimation of the noise power by computing the data-template residual at the maximum-likelihood parameters. We also propose an

iterative approach to using the GPR method, where the training set and noise estimate are refined through (a single) repetition. Finally, in Sec. III G, we implement all of the above improvements to perform inference on an example EMRI signal with simulated LISA noise, and to demonstrate the viability of the GPR marginalized-likelihood method for more realistic GW applications.

ACKNOWLEDGMENTS

M. L. would like to express his gratitude to Professor Yiming Hu and Professor Jiandong Zhang for their insightful discussions and valuable feedback on our research. M. L. is also grateful to the NUS Research Scholarship for its financial support. A. J. K. C. thanks Christopher Moore for helpful comments on the manuscript, and acknowledges previous support from the NASA LISA Preparatory Science Grant No. 20-LPS20-0005.

-
- [1] R. Abbott, T. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. Adhikari, V. Adya, C. Affeldt, D. Agarwal *et al.*, [arXiv:2111.03606](#).
- [2] R. Abbott *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), *Phys. Rev. X* **13**, 011048 (2023).
- [3] P. Amaro-Seoane *et al.* (LISA Collaboration), [arXiv:1702.00786](#).
- [4] J. Mei, Y.-Z. Bai, J. Bao, E. Barausse, L. Cai, E. Canuto, B. Cao, W.-M. Chen, Y. Chen, Y.-W. Ding *et al.*, *Prog. Theor. Exp. Phys.* **2021**, 05A107 (2021).
- [5] J. Luo, L.-S. Chen, H.-Z. Duan, Y.-G. Gong, S. Hu, J. Ji, Q. Liu, J. Mei, V. Milyukov, M. Sazhin *et al.*, *Classical Quantum Gravity* **33**, 035010 (2016).
- [6] W.-H. Ruan, C. Liu, Z.-K. Guo, Y.-L. Wu, and R.-G. Cai, *Nat. Astron.* **4**, 108 (2020).
- [7] S.-J. Huang, Y.-M. Hu, V. Korol, P.-C. Li, Z.-C. Liang, Y. Lu, H.-T. Wang, S. Yu, and J. Mei *et al.*, *Phys. Rev. D* **102**, 063021 (2020).
- [8] V. Korol, E. M. Rossi, P. J. Groot, G. Nelemans, S. Toonen, and A. G. A. Brown, *Mon. Not. R. Astron. Soc.* **470**, 1894 (2017).
- [9] H.-T. Wang, Z. Jiang, A. Sesana, E. Barausse, S.-J. Huang, Y.-F. Wang, W.-F. Feng, Y. Wang, Y.-M. Hu, J. Mei *et al.*, *Phys. Rev. D* **100**, 043003 (2019).
- [10] A. Klein *et al.*, *Phys. Rev. D* **93**, 024003 (2016).
- [11] S. Liu, Y.-M. Hu, J.-d. Zhang, and J. Mei *et al.*, *Phys. Rev. D* **101**, 103027 (2020).
- [12] A. Sesana, *Phys. Rev. Lett.* **116**, 231102 (2016).
- [13] S. Babak, J. Gair, A. Sesana, E. Barausse, C. F. Sopuerta, C. P. Berry, E. Berti, P. Amaro-Seoane, A. Petiteau, and A. Klein, *Phys. Rev. D* **95**, 103012 (2017).
- [14] H.-M. Fan, Y.-M. Hu, E. Barausse, A. Sesana, J.-d. Zhang, X. Zhang, T.-G. Zi, and J. Mei, *Phys. Rev. D* **102**, 063016 (2020).
- [15] S. Babak, J. Gair, A. Sesana, E. Barausse, C. F. Sopuerta, C. P. L. Berry, E. Berti, P. Amaro-Seoane, A. Petiteau, and A. Klein, *Phys. Rev. D* **95**, 103012 (2017).
- [16] C. P. L. Berry, S. A. Hughes, C. F. Sopuerta, A. J. K. Chua, A. Heffernan, K. Holley-Bockelmann, D. P. Mihaylov, M. C. Miller, and A. Sesana, [arXiv:1903.03686](#).
- [17] J. R. Gair, C. Tang, and M. Volonteri, *Phys. Rev. D* **81**, 104014 (2010).
- [18] P. Canitrot, *Phys. Rev. D* **63**, 082005 (2001).
- [19] C. Cutler and M. Vallisneri, *Phys. Rev. D* **76**, 104018 (2007).
- [20] A. H. Mroue *et al.*, *Phys. Rev. Lett.* **111**, 241104 (2013).
- [21] J. Aasi *et al.* (LIGO Scientific, Virgo, and NINJA-2 Collaborations), *Classical Quantum Gravity* **31**, 115004 (2014).
- [22] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms* (Cambridge University Press, Cambridge, England, 2003).
- [23] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (MIT Press, Cambridge, MA, 2006).
- [24] C. J. Moore and J. R. Gair, *Phys. Rev. Lett.* **113**, 251101 (2014).
- [25] C. J. Moore, C. P. L. Berry, A. J. K. Chua, and J. R. Gair, *Phys. Rev. D* **93**, 064001 (2016).
- [26] A. J. K. Chua, *J. Phys. Conf. Ser.* **716**, 012028 (2016).
- [27] C. J. Moore, A. J. K. Chua, C. P. L. Berry, and J. R. Gair, *R. Soc. Open Sci.* **3**, 160125 (2016).
- [28] A. J. K. Chua, N. Korsakova, C. J. Moore, J. R. Gair, and S. Babak, *Phys. Rev. D* **101**, 044027 (2020).
- [29] M. L. Katz, A. J. K. Chua, L. Speri, N. Warburton, and S. A. Hughes, *Phys. Rev. D* **104**, 064047 (2021).
- [30] A. Pound, B. Wardell, N. Warburton, and J. Miller, *Phys. Rev. Lett.* **124**, 021101 (2020).

- [31] N. Warburton, A. Pound, B. Wardell, J. Miller, and L. Durkan, *Phys. Rev. Lett.* **127**, 151102 (2021).
- [32] B. Wardell, A. Pound, N. Warburton, J. Miller, L. Durkan, and A. Le Tiec, *Phys. Rev. Lett.* **130**, 241402 (2023).
- [33] A. J. K. Chua, M. L. Katz, N. Warburton, and S. A. Hughes, *Phys. Rev. Lett.* **126**, 051102 (2021).
- [34] L. Speri, M. L. Katz, A. J. K. Chua, S. A. Hughes, N. Warburton, J. E. Thompson, C. E. A. Chapman-Bird, and J. R. Gair, [arXiv:2307.12585](https://arxiv.org/abs/2307.12585).
- [35] C. Cutler and E. E. Flanagan, *Phys. Rev. D* **49**, 2658 (1994).
- [36] C. Cutler and M. Vallisneri, *Phys. Rev. D* **76**, 104018 (2007).
- [37] A. J. K. Chua, C. J. Moore, and J. R. Gair, *Phys. Rev. D* **96**, 044005 (2017).
- [38] R. Fujita and M. Shibata, *Phys. Rev. D* **102**, 064005 (2020).
- [39] S. Isoyama, R. Fujita, A. J. K. Chua, H. Nakano, A. Pound, and N. Sago, *Phys. Rev. Lett.* **128**, 231101 (2022).
- [40] N. Sago and R. Fujita, *Prog. Theor. Exp. Phys.* **2015**, 073E03 (2015).
- [41] C. Cutler, *Phys. Rev. D* **57**, 7089 (1998).
- [42] M. Tinto and J. W. Armstrong, *Phys. Rev. D* **59**, 102003 (1999).
- [43] M. Tinto, F. B. Estabrook, and J. W. Armstrong, *Phys. Rev. D* **65**, 082003 (2002).
- [44] A. Krolak, M. Tinto, and M. Vallisneri, *Phys. Rev. D* **70**, 022003 (2004); **76**, 069901(E) (2007).
- [45] M. C. Shewry and H. P. Wynn, *J. Appl. Stat.* **14**, 165 (1987).
- [46] C. Sheng, M. H. Y. Tan, and L. Zou, *J. Appl. Stat.* **48**, 2152 (2021).
- [47] M. D. McKay, R. J. Beckman, and W. J. Conover, *Technometrics* **42**, 55 (2000).
- [48] M. D. Morris and T. J. Mitchell, *J. Stat. Plann. Infer.* **43**, 381 (1995).
- [49] M. E. Johnson, L. M. Moore, and D. Ylvisaker, *J. Stat. Plann. Infer.* **26**, 131 (1990).
- [50] M. A. Bouhlel, J. T. Hwang, N. Bartoli, R. Lafage, J. Morlier, and J. R. R. A. Martins, *Adv. Eng. Softw.* **135**, 102662 (2019).
- [51] M. R. Hestenes, E. Stiefel *et al.*, *J. Res. Natl. Bur. Stand.* **49**, 409 (1952).
- [52] T. Robson, N. J. Cornish, and C. Liu, *Classical Quantum Gravity* **36**, 105011 (2019).
- [53] W. K. Hastings, *Biometrika* **57**, 97 (1970).
- [54] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).