

One-stop function for gravitational-wave detection, identification, and inference

Alvin J. K. Chua^{*}

*Department of Physics, National University of Singapore 117551, Singapore;
Department of Mathematics, National University of Singapore 119076, Singapore
and Theoretical Astrophysics Group, California Institute of Technology,
Pasadena, California 91125, USA*

 (Received 31 August 2022; accepted 27 October 2022; published 28 November 2022)

I define here a novel function on a modeled space of gravitational-wave signals, before studying its properties as a statistic for detection, as an objective function for identification, and as an effective likelihood function for inference. The main motivation behind this work is the open data-analysis problem for signals from extreme-mass-ratio inspirals, which is severely hindered by the presence of strong nonlocal parameter degeneracy in the signal space. I demonstrate the utility of the proposed function for the analysis of such signals, and suggest various possible directions for future research.

DOI: [10.1103/PhysRevD.106.104051](https://doi.org/10.1103/PhysRevD.106.104051)

I. INTRODUCTION

The scientific inverse problem in gravitational-wave (GW) astronomy comprises three distinct data-analysis procedures (in addition of course to the vital tasks of noise characterization and forward source modeling, both of which we will assume can be achieved at full accuracy for present purposes). In the order they are performed, these procedures are:

- (i) Detection: Establishing the statistically significant presence of a GW signal in noisy detector data;
- (ii) Identification: Mapping the detected signal (sufficiently) accurately to the source parameters of a (sufficiently) representative forward model;
- (iii) Inference: Estimating the Bayesian posterior probability of the actual source parameters.

Detection can be performed with or without models, while identification requires at least an approximate model, and inference relies on a highly accurate one. The combination of detection and identification is commonly referred to as “search.” For *almost* all classes of GW source, identification is unnecessary when inference in the full model space is feasible, and/or guaranteed when accurate models are straightforward to invert (hence its complete neglect in much of the literature).

In the case of the extreme-mass-ratio inspirals (EMRIs) that will be observed by the near-future ESA–NASA mission LISA [1–4], the traditional dichotomy of detection and inference (or “parameter estimation”) breaks down. Indeed, the overall EMRI inverse problem remains an open one, with well-documented obstacles arising from the

theoretical challenges in achieving the required accuracy and extensiveness for models [5,6]; computational limitations in attaining sufficient model efficiency [7,8]; as well as the inherent difficulty in exploring the voluminous EMRI signal space, due to *strong and nonlocal parameter degeneracy*. This last feature was recently characterized qualitatively by Chua and Cutler [9], and is the primary hindrance to all three procedures of detection (uncertain criteria for reliable candidates), identification (multiple spurious candidates across signal space), and inference (impractical without highly localized priors) within the EMRI data-analysis problem. Past work on EMRI search [10–14] has sought to provide practical solutions, but has not shown that identification in particular can be accomplished both reliably and efficiently.

For a fresh take on the problem, I propose in this manuscript a “one-stop” function with various favorable geometrical and statistical properties for EMRI detection, identification and inference. These properties are analytically and numerically characterized to an extent that should in my opinion leave little doubt over the utility of the function in realistic applications; stress tests on full-scale EMRI simulations are consequently left for future studies. The function takes the form

$$f(\theta) := X(\theta) \exp\left(-\frac{1}{2}\beta(\theta)\chi^2(\theta)\right). \quad (1)$$

Here, X is the standard matched-filtering detection statistic for a single signal template [Eq. (5)]; $\beta > 0$ is a temperaturelike calibration factor that is *fully specified by the template model* [Eq. (17)]; and χ^2 reduces for actual signals to a chi-squared statistic with M degrees of freedom

^{*}alvincjk@nus.edu.sg

[Eq. (14)], arising from a suitable decomposition of the template model into $M > 1$ modes. Equation (1) is then to be used in a stochastic sampling algorithm, where the target probability density is $\propto \exp f$ (and so its logarithm, which is the commonly supplied quantity, is $f + \text{const}$). It has no tuning parameters.

The manuscript also touches on a larger second thesis: that most classes of coherent (phase-matching) statistics employed in stochastic GW searches are actually suboptimal due to *uncontrolled variations* of the search statistic over the model space, caused by nonlocal signal correlations as well as the manifestation of detector noise. Such variations can generally be eliminated in f for any specific application of coherent search, through the choice of χ^2 (the mode decomposition) and β . While these are tailored here to the qualitative nature of the EMRI signal space, they might also be defined for example to suppress the impact of transient noise artifacts, or to aid future searches for long-duration signals from precessing and eccentric comparable-mass binaries (where nonlocal degeneracy could likewise be an issue).

II. DERIVATION

In this section, we will expand on the derivation of Eq. (1) in some detail. Our starting point is the generic decomposition of a GW-signal template $h(\theta)$ as

$$h(\theta) = \sum_{m=1}^M h_m(\theta) + \epsilon(\theta), \quad (2)$$

with the following assumptions, valid for all m, m' :

$$\langle h_m | h_m \rangle > 1, \quad |\langle h_m | h_{m'} \rangle| \ll 1, \quad \langle \epsilon | \epsilon \rangle \ll 1, \quad (3)$$

where $\langle \cdot | \cdot \rangle$ is the usual detector-noise-weighted inner product on the space of fixed-length time series [15]. We will use the term ‘‘modes’’ to refer to any set $\{h_m\}$ satisfying (2) and (3), rather than the familiar angular modes $\{h_{lm}\}$ of gravitational radiation. (In fact, $\{h_{lm}\}$ generally does not satisfy the first two conditions of (3).)

The decomposition (2) and (3) can still be achieved phenomenologically for comparable-mass binary mergers by a simple partition of their signals in the frequency domain [16], and thus the proposed function (1) might be beneficial (but not nearly as crucial) for the analysis of such sources. Equations (2) and (3) might also be satisfied by the near-monochromatic signals from quasicircular early-inspiral binaries, through a partition in the time domain. A more physically motivated decomposition can be obtained for generically inclined and eccentric EMRIs into a Kerr black hole, where the signal power is spread across various harmonics of the three fundamental frequencies; indeed, the canonical modeling approach relies on an angular and

frequency-based decomposition $h = \sum h_{lmkn}$ in the first place [17].

Let us now quickly revisit various familiar quantities and identities in GW data analysis, before introducing several new ones. The optimal signal-to-noise ratio (SNR) for a template h is its norm with respect to $\langle \cdot | \cdot \rangle$:

$$\rho(\theta) := \sqrt{\langle h(\theta) | h(\theta) \rangle}. \quad (4)$$

We will denote the time-series data from the detector as x , and work primarily with two scenarios (hypotheses): the null H_0 , where x equals the (assumed stationary and Gaussian) detector noise n ; and H_1 , where $x = h^* + n$ for a single EMRI signal h^* . Further assume that h^* lies in the signal space described by the template model $h(\theta)$, such that $h^* = h(\theta^*)$ for some θ^* . The detection SNR for a template h is the scalar projection of x on h , and is a statistic of the data (so we will henceforth refer to it as the standard detection statistic). It is defined as

$$X(\theta) := \langle x | \hat{h}(\theta) \rangle := \frac{\langle x | h(\theta) \rangle}{\rho(\theta)}. \quad (5)$$

Recall that for all time series a, b , the noise-weighted inner product satisfies the identities

$$\mathbb{E}[\langle n | a \rangle] = 0, \quad \mathbb{E}[\langle n | a \rangle \langle n | b \rangle] = \langle a | b \rangle. \quad (6)$$

Thus at the template parameters θ^* (with corresponding optimal SNR ρ^*), we have

$$X | H_0 \sim \mathcal{N}(0, 1), \quad X | H_1 \sim \mathcal{N}(\rho^*, 1). \quad (7)$$

Finally, we have the standard Bayesian (log-)likelihood function of the source parameters θ , given H_1 :

$$\begin{aligned} \ln L(\theta) &:= -\frac{1}{2} \langle x - h(\theta) | x - h(\theta) \rangle \\ &= \langle x | h \rangle - \frac{1}{2} (\langle h | h \rangle + \langle x | x \rangle). \end{aligned} \quad (8)$$

There is a clearly a close relationship between X and L ; we can make this more intuitive by observing that the term $\langle x | x \rangle$ is constant over parameter space, while we also have $\partial_\theta \langle h | h \rangle \ll \partial_\theta \langle x | h \rangle$. Thus $\exp(X\rho) = \exp \langle x | h \rangle$ is an excellent local approximation to L as a density function (i.e., modulo a normalization factor).

We may now define analogous vector-valued versions of ρ and X using the mode decomposition (2), and denote these using boldface by $\boldsymbol{\rho}$ and \mathbf{X} . These M -vectors are given component-wise by

$$[\boldsymbol{\rho}(\theta)]_m := \rho_m(\theta) := \sqrt{\langle h_m(\theta) | h_m(\theta) \rangle}, \quad (9)$$

$$[\mathbf{X}(\theta)]_m := X_m(\theta) := \langle x | \hat{h}_m(\theta) \rangle := \frac{\langle x | h_m(\theta) \rangle}{\rho_m(\theta)}. \quad (10)$$

It is straightforward to see that

$$\mathbf{X} \cdot \boldsymbol{\rho} = \langle x | h \rangle = X\rho \quad (11)$$

and, from the second assumption in Eq. (3), that

$$|\boldsymbol{\rho}|^2 = \boldsymbol{\rho} \cdot \boldsymbol{\rho} = \langle h | h \rangle = \rho^2, \quad (12)$$

but note that $|\mathbf{X}|^2 = \mathbf{X} \cdot \mathbf{X} \neq \langle x | x \rangle$.

From the noise identities (6) [and the second assumption in Eq. (3)], we have for θ^* and all m :

$$X_m | H_0 \sim \mathcal{N}(0, 1), \quad X_m | H_1 \sim \mathcal{N}(\rho_m^*, 1), \quad (13)$$

where the ρ_m^* sum to ρ in quadrature from Eq. (12). Thus the chi-squaredlike statistic

$$\chi^2(\theta) := |\mathbf{X}(\theta) - \boldsymbol{\rho}(\theta)|^2 \quad (14)$$

has a chi-squared distribution with M degrees of freedom, but only given H_1 and when evaluated at the point θ^* .

The appearance of χ^2 in the exponential factor of f suppresses secondary peaks arising from nonlocal parameter degeneracy, which lifts the main barrier to EMRI search and inference. At this point, the observant reader might note that the exponential amplification of *any* discrepancy between the signal and the secondary template also fulfills the same purpose. Crucially, however, the additional information provided by the effective decomposition of the detection statistic X into $M > 1$ modes allows secondary suppression to be achieved without a severe impact on the overall detection sensitivity of f at θ^* . We will return to this point shortly in Sec. II A.

Let us now turn to the remaining undefined quantity in Eq. (1), which is the calibration factor β . Without loss of generality, we may sort the decomposition in Eq. (2) by imposing $\langle h_m | h_m \rangle \leq \langle h_{m+1} | h_{m+1} \rangle$, such that the dominant mode is h_M (which is often h_{2200} , in large regions of the signal space). One of the key findings in Chua and Cutler [9] is that the strongest secondary peaks in the likelihood (or detection statistic) over the model parameter space are overwhelmingly caused by the dominant mode of a nonlocal template being phase-aligned with the dominant mode of the actual signal, *without significant contribution from any alignment of the other modes*.

We may use this observation to calibrate the secondary suppression of Eq. (1) with $n = 0$ against the noise properties of $X | H_0$; this can be performed using only information about the signal space. For each θ , let us consider the corresponding putative signal $h(\theta)$, and then require that $f^2 = \text{Var}(X) = 1$ for any putative template with a dominant mode that perfectly matches that of

$h(\theta)$, but no other matched modes. Such templates are not explicitly required, of course, since they satisfy by definition

$$X_m = \rho_m(\theta) \delta_{mM}, \quad (15)$$

where δ is the Kronecker delta function. For these templates, it follows that

$$f = \rho_M(\theta) \exp\left(-\frac{1}{2}\beta(\theta)(\rho^2(\theta) - \rho_M^2(\theta))\right) = 1. \quad (16)$$

Thus we arrive at

$$\beta(\theta) := \frac{2 \ln(\alpha(\theta)\rho(\theta))}{(1 - \alpha(\theta)^2)\rho(\theta)^2}, \quad (17)$$

where we have defined $\alpha := \rho_M/\rho$ as the fractional optimal SNR of the dominant mode relative to the full template at each parameter point (and so α^2 is the fractional power). Note that $\beta(\alpha)$ is monotonically increasing; also,

$$1 < M < \rho^2, \quad \frac{1}{\sqrt{M}} < \alpha < \frac{\sqrt{\rho^2 - M + 1}}{\rho} \quad (18)$$

from the first assumption in Eq. (3), so β is well defined:

$$\frac{(\ln \rho^2 - \ln M)M}{\rho^2(M - 1)} < \beta < \frac{\ln(\rho^2 - M + 1)}{M - 1}. \quad (19)$$

Finally, this particular definition of β allows Eq. (1) to be written alternatively as

$$f = X(\alpha\rho)^{-\chi^2/((1-\alpha^2)\rho^2)}, \quad (20)$$

which is useful for direct evaluation—but we shall keep working with the exponential form for analytical convenience. The end result of the calibration is, essentially, that any secondary peaks in $f | H_1$ arising from nonlocal degeneracy will have a maximum height comparable to the typical variation of $X | H_0$ due to detector noise.

A. Relationship to other functions

It is useful to construe the proposed function f as an approximation to a more easily interpreted function

$$f'(\theta) := X(\theta) \left(1 - \gamma\left(\frac{M}{2}, \frac{\beta'(\theta)\chi^2(\theta)}{2}\right)\right), \quad (21)$$

where γ denotes the regularized lower incomplete gamma function (i.e., the cumulative distribution function of $\beta'\chi^2$ if it were chi-squared-distributed). The approximation of f' by f is exact in the case $M = 2$ and $\beta' = \beta$. For $M > 2$, the approximation is still reasonably preserved by the

calibration factor β' , which is now $\neq \beta$ but similarly determined by requiring $f' = 1$ at the location of a pure dominant-mode secondary in the case of $n = 0$:

$$\beta' = \frac{2}{(1 - \alpha^2)\rho^2} \gamma^{-1} \left(\frac{M}{2}, 1 - \frac{1}{\alpha\rho} \right). \quad (22)$$

The flatter profile in the vicinity of θ^* of the $1 - \gamma$ factor relative to the exponential factor in Eq. (1) means that f' is locally a closer approximation to X than f , which could be beneficial in practical terms. Nevertheless, the function f' is more difficult than f to characterize analytically, and is thus left for future investigation.

When β' is simply set to unity across the signal space, f' evokes the chi-squared-weighted SNR seen in ground-based pipelines (as introduced in [18]). There, however, X is weighted by an empirically determined function of χ^2 to suppress nonstationary/non-Gaussian noise, and the product of the two has not been fundamentally characterized in the manner of Sec. III (to the best of my knowledge). The $\beta' = 1$ case is also equivalent to the modified $\langle x|h \rangle$ term in the first of the two “veto likelihoods” proposed in Chua and Cutler [9], modulo a factor of ρ . The functions in that paper also fulfill the specific aim of secondary suppression, but in a more *ad hoc* way and, crucially, as an adjustment to $\ln L$ rather than X .

There exists another interesting interpretation of f as the product of X and L' , where the (scaled) logarithm of the latter is written in a deliberately evocative way as

$$\frac{\ln L'}{\beta} := \langle x|h \rangle - \frac{1}{2} \left(\langle h|h \rangle + \sum_m \langle x|\hat{h}_m \rangle \langle \hat{h}_m|x \rangle \right). \quad (23)$$

The function L' bears obvious structural similarities to L , with the key difference being that the squared norm (with respect to $\langle \cdot | \cdot \rangle$) of x is replaced with the squared norm (with respect to the Euclidean inner product) of its projection onto the mode basis describing each template. One may then wonder whether the function XL , or even XL' with $M = 1$, might provide a similar suppression of secondaries with suitable calibration. However, there is no way to define a sufficiently discriminative calibration without mode information—any attempt leads to an over-suppression of templates in the local vicinity of θ^* , which is further amplified in the presence of noise and destroys the sensitivity of the template at θ^* itself.

III. PROPERTIES

To study the various geometrical and statistical properties of the function f for EMRI data analysis, we will make use of the AAK template model for generically inclined and eccentric Kerr EMRIs [19,20], along with the long-wavelength LISA response [21]. We keep to the slightly older implementation employed in Chua and Cutler [9], rather than the latest version introduced in [8] (to avoid

redoing the computationally tedious search for secondaries in the space of a new model). As in [9], our analysis is restricted to six intrinsic parameters: the two component masses, the spin of the primary mass, and the three initial conditions for the orbit/frequencies. These are commonly parametrized as

$$\theta := \left(\lg \frac{\mu}{M_\odot}, \lg \frac{M}{M_\odot}, \frac{a}{M}, \frac{p_0}{M}, e_0, \cos i_0 \right), \quad (24)$$

where $(\mu, M \gg \mu)$ are the masses, a is the usual spin length scale describing Kerr spacetime, and (p_0, e_0, i_0) are the initial quasi-Keplerian semi-latus rectum, eccentricity and inclination of the orbit.

The signal injection and “secondary B” from [9] are chosen here as a representative injection and secondary template, respectively; their associated parameters are

$$\theta^* = (1, 6, 0.5, 9.5, 0.2, 0.866), \quad (25)$$

$$\theta^S = \theta^* + (3, -12, -54, 170, 1, 64) \times 10^{-3}. \quad (26)$$

While the secondary parameters might not look particularly nonlocal to the uninitiated reader at first glance, they lie *well beyond* any realistic Bayesian credible region for the inferred parameters. At a typical SNR of 20, the 1-sigma values for these parameters (relative to the maximum *a posteriori* estimate) are

$$\pm(0.07, 0.2, 0.5, 3, 0.06, 0.8) \times 10^{-3}. \quad (27)$$

There are however some key changes between the setup here and in [9]. For the present analysis, we consider only full templates sampled at 0.1 Hz, and not the phase-trajectory “templates” adopted in that work. White noise and the Euclidean inner product on time series are used here for convenience, and without much loss of generality—this simplification has negligible impact on the correlation structure of signal space at the considered scales, while the properties of any statistics derived from the noise-weighted inner product are preserved. More crucially, the exact definition of modes is slightly different in the two studies. In [9], only $M = 4$ radial modes with $(l, m, k) = (2, 2, 0)$ and $-1 \leq n \leq 2$ are considered; here, we still use $M = 4$, $l = 2$ and $-1 \leq n \leq 2$, but with an implicit sum over (m, k) (i.e., the modes now include their sidebands from Lense–Thirring precession, and essentially form the partially decomposed quadrupolar harmonic basis introduced by Barack and Cutler [22]). Finally, the signals we examine here are shortened from a duration of two years to six months, but renormalized accordingly to retain an injection SNR of 20. This does not change the location of strong secondaries, at least not significantly (see Sec. IV B 1 in [9] to understand why).

A. As a detection statistic

Let us now examine the statistical properties of $f(\theta^*)$ under the two hypotheses $H_0: x = n$ and $H_1: x = h^* + n$. Closed-form descriptions of the full probability distributions are not straightforward to obtain, but their first couple of moments are comfortably within reach:

$$E[f|H_0] = \frac{\beta\rho}{(\beta+1)^{(M+2)/2}} \exp\left(-\frac{\beta\rho^2}{2(\beta+1)}\right), \quad (28)$$

$$E[f^2|H_0] = \frac{4\beta^2\rho^2 + 2\beta + 1}{(2\beta+1)^{(M+4)/2}} \exp\left(-\frac{\beta\rho^2}{2\beta+1}\right), \quad (29)$$

$$E[f|H_1] = \frac{\rho}{(\beta+1)^{M/2}}, \quad (30)$$

$$E[f^2|H_1] = \frac{(2\beta+1)\rho^2 + 1}{(2\beta+1)^{(M+2)/2}}. \quad (31)$$

The full distributions of $f|H_0$ and $f|H_1$ are strongly non-Gaussian, and so the moments (28)–(31) are only loosely indicative of their location and scale. Nevertheless, the usual means $E[f]$ and 1-sigma values $E[f] \pm (E[f^2] - E[f]^2)^{1/2}$ are convenient quantities for examining how these distributions depend on α , M and ρ^* . In Fig. 1, the mean and spread of $f|H_0$ and $f|H_1$ are plotted as a function of α for the representative values of $M = 4$ and $\rho^* = 20$. Also included are curves for an increased number of modes ($M = 30$ is approximately the minimum number to describe EMRIs of all eccentricities in the Barack–Cutler harmonic basis [22]), as well as for a slightly higher injection SNR of 30.

For all considered M and ρ^* , the mean and spread of $f|H_0$ are $\ll 1$ for most values of α —i.e., the statistic f is highly insensitive to noise. Its sensitivity to an actual signal is also reduced from that of X , although this reduction is limited for $M = 4$ and $\alpha \lesssim 0.8$ (since $f|H_1$ retains a mean of $\approx \rho^*$ with a spread of ≈ 1 in that regime). As α approaches its maximal value in Eq. (18), the sensitivity of f vanishes, which diminishes its utility for near-equatorial Kerr EMRIs with eccentricities that are $\ll 1$ (since $h \approx h_M = h_{2200}$). This is a narrow region of signal space, however, and higher l -modes might still carry enough power for α to be suppressed in parts of the region (if sufficient SNR is accumulated close to plunge).

Even with the Barack–Cutler mode decomposition, a fairly large number of modes must be included in Eq. (2) for a *global* description of the signal space [22]. This is detrimental to the sensitivity of f (see red curve in Fig. 1)—less so for high-eccentricity sources, where α is lower due to the larger spread of power across the modes, but certainly for low-eccentricity ones, which would be better analyzed with fewer modes. However, as M is a discrete parameter, we will prefer to avoid varying it across signal space. It is

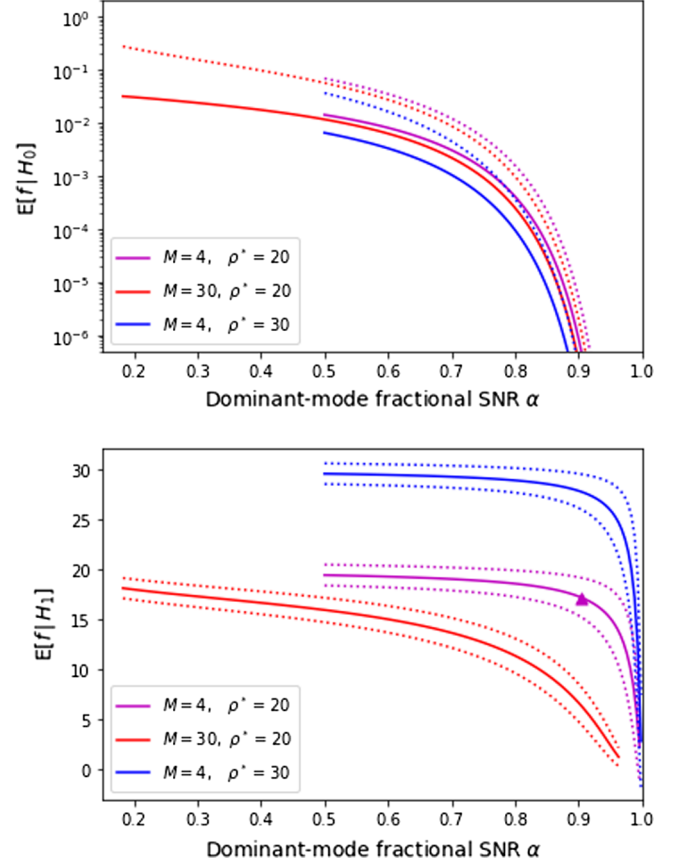


FIG. 1. Top: dependence of the means (solid lines) and 1-sigma values (dotted lines) of $f(\theta^*)|H_0$ on α , for various M and ρ^* . Note that the vertical axis is in logarithmic scale. Bottom: dependence of the means and 1-sigma values of $f(\theta^*)|H_1$ on α . The triangle corresponds to the representative injection considered in this work, for which $\alpha \approx 0.9$.

clear from Eq. (30) and Fig. 1 that the ideal decomposition for each signal in the space is with both α and M as small as possible. This might be accomplished for fixed M through a θ -dependent definition of *the decomposition itself*, e.g., relaxing the third assumption in Eq. (3) by just choosing the M strongest modes for each template, or even treating ϵ (the sum of the remaining modes) as a single mode if it satisfies the first two assumptions and is not dominant. Such variants ensure that f still changes smoothly over the signal space; we will leave their study to future global analyses.

To characterize the discriminative power of f as a detection statistic for the representative signal injection considered in this work, we may empirically estimate $p(f|H_0)$ and $p(f|H_1)$ by generating 10^4 realizations of each statistic, then fitting the samples to a large family of common distributions with a Kolmogorov–Smirnov test. The empirical distribution for $f|H_0$ admits no satisfactory fits, but is observed to satisfy $P(f > 0|H_0) \approx 1/2$ to a very good approximation. Since we are concerned only with the positive realizations of $f|H_0$ (denoted by $f^+|H_0$) when computing false-alarm thresholds, we may fit only these

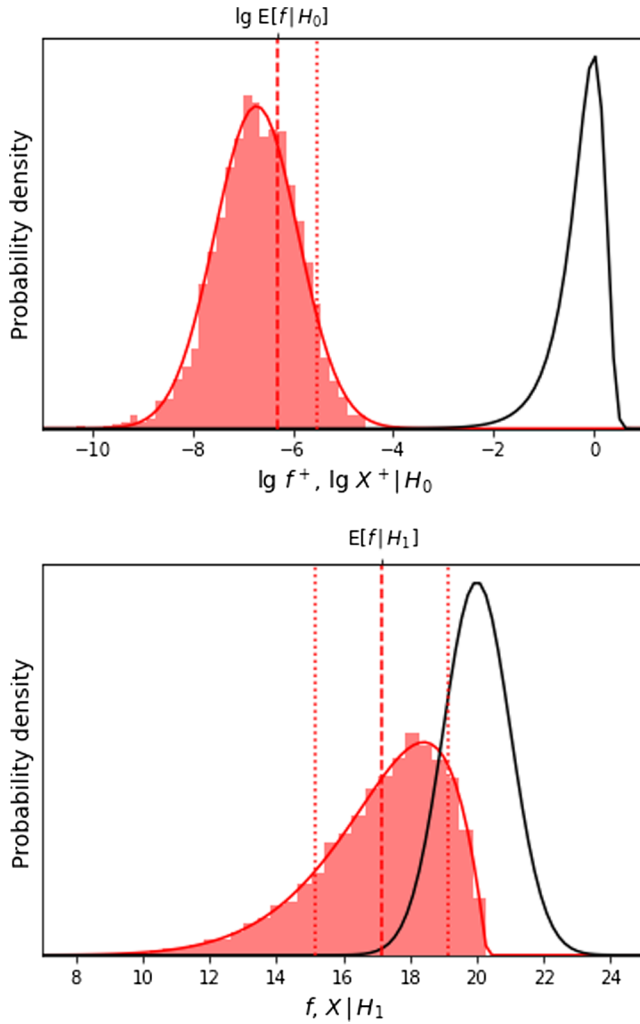


FIG. 2. Top: probability distributions of $\lg f^+(\theta^*)|H_0$ (red; empirical and normal fit) and $\lg X^+(\theta^*)|H_0$ (black) for the representative injection considered in this work. Bottom: probability distributions of $f(\theta^*)|H_1$ (red; empirical and generalized-beta fit) and $X(\theta^*)|H_1$ (black). In both plots, the exact means (dashed lines) and 1-sigma values (dotted lines) of $f(\theta^*)$ are included for reference.

and simply reduce the fitted probability density by a factor of two. The distribution of $f^+|H_0$ is best fit by a log-normal distribution with a Kolmogorov–Smirnov statistic $D_{10^4} \approx 0.01$, while $f|H_1$ obeys a generalized beta distribution with $D_{10^4} \approx 0.006$ (see Fig. 2).

With the fitted probability densities $p(f^+|H_0)$ and $p(f|H_1)$ in hand, we may then estimate the false-alarm probability for f given some fixed detection probability. To achieve a minimum detection probability

$$P_D := \int_{f_T}^{\infty} df p(f|H_1) = 1 - 10^{-3}, \quad (32)$$

the maximum threshold value for f is $f_T \approx 9$, and the corresponding false-alarm probability is

$$P_F := \int_{f_T}^{\infty} df^+ \frac{1}{2} p(f^+|H_0) \sim 10^{-21}. \quad (33)$$

An analogous calculation for the standard detection statistic X yields a threshold $X_T \approx 17$ and $P_F \sim 10^{-64}$. Thus f is *formally* less discriminative than X due to both its reduced sensitivity and the heavier tail of its (fitted) null distribution, although the practical difference between the two is negligible at an SNR of 20 (and this representative choice of θ^*).

Across the full signal space, nonlocal degeneracy greatly complicates any attempt at global analysis, and diminishes the validity of idealized treatments, e.g., highly conservative estimates of the required detection-statistic threshold at fixed false-alarm probability for a putative bank of uncorrelated templates $\{h(\theta_i)\}$ spanning the space [20,23]. Furthermore, the joint distributions of the corresponding statistics $\{f_i\}$ under both hypotheses depend strongly on the mode structure across signal space, as opposed to those for $\{X_i\}$. The single-template analysis from above would seem to indicate that $\{f_i\}$ is more prone to Type-II errors than $\{X_i\}$ for a fixed false-alarm probability; however, it also completely fails to address global issues such as the increased Type-I error rate of $\{X_i\}$ over $\{f_i\}$ when detecting the exact set of signals present in data, since it cannot account for false alarms due to secondaries of actual signals.

B. As an objective function

The identification aspect of search is in essence an optimization problem—specifically, the global maximization of some objective function that describes how well the data matches each point in the space of a template model. Pure optimization techniques are highly efficient when maximizing concave(-down) or even log-concave functions, while optimization through sampling is more suitable for functions that are only concave or log-concave beyond a bounded region [24]. Sampling, or at least stochastic optimization, would seem to be the only viable option for nontextbook objective functions such as X and L . Although these functions are log-concave in a highly localized region, their large-scale structure is not (unless a globally log-concave prior is specified, but this rather defeats the purpose of search). Furthermore, they suffer not only from intrinsic variations due to nonlocal signal correlations (with EMRIs being the most extreme manifestation of this), but also from noise variations that are locally correlated on the same length scales as the points in the template model.

While the function f is derived with secondary suppression as the primary motivation, it has the beneficial side effect of significantly suppressing noise as well; this fact turns out to be quite useful for the purposes of identification. In Fig. 3, the values of $f|H_1$ and $X|H_1$ are compared along an (extended) connecting line in parameter space between the injection parameters θ^* and the considered




FIG. 3. Top: extended connections of X , S_6 , S_{12} and f between injection and secondary parameters, without noise ($x = h^*$). Also included are an annealed standard detection statistic X/ρ^* , as well as f for a different signal space (with a truncated duration of three months from the reference initial time). Bottom: as above, but with noise ($x = h^* + n$).

secondary parameters θ^S . Following Chua and Cutler [9], we will refer to the restriction of each function to this one-dimensional domain as a *connection*. Not shown in the figure is $\ln L$, which is effectively $X\rho^*$ with a slight gradient in the baseline due to changes in $\rho(\theta)$ (these become significant on global length scales, and so the likelihood is ill suited to search). Alternatively, $\exp X$ is locally an annealed version of L with temperature ρ^* .

Also included in the comparison is a semicoherent statistic S_N where both data and template are transformed into the time–frequency domain (with a time resolution of N), and the statistic is the sum of N *less informative* statistics defined on the individual time segments. This concept is most simply realized through a short-time discrete Fourier transform of both time series, with a partition into N similar-length segments (indexed by i) and no window. We may define a semi-coherent inner product between data and template as

$$\langle x|h \rangle_N := \sum_{i=1}^N \max_{\text{ext.}} \langle x_i|h_i \rangle, \quad (34)$$

where the maximization is over a subset of the extrinsic degrees of freedom. These are not the extrinsic parameters *per se*, and can bleed into the intrinsic parameters in the standard EMRI parametrization. Note that $\langle x|h \rangle_1 \neq \langle x|h \rangle$; also, for white noise (or whitened time series with the Euclidean inner product), each $\langle x|h \rangle_N$ reduces exactly to $\langle x|h \rangle$ in the absence of maximization.

For present purposes, we will restrict to analytically maximizing over an overall time shift [25], but not an overall phase shift [26] for all modes (which is less trivial to perform analytically in the case of EMRIs). Our semicoherent statistic itself is then defined as

$$S_N := \frac{\langle x|h \rangle_N}{\sqrt{\langle h|h \rangle_N}} = \frac{\langle x|h \rangle_N}{\rho} = \langle x|\hat{h} \rangle_N. \quad (35)$$

Maximizing over some extrinsic degrees of freedom naturally simplifies the global structure of S_N , at the cost of losing information about those degrees of freedom. This manifests as a significantly broadened peak for S_N around θ^* , where the broadening increases with N (see Fig. 3). More generally, semicoherent filtering for EMRI search [27] is an oft-discussed but as yet unactualized paradigm, since it provides an effective means of constructing objective functions with similar properties to S_N . It is also highly suited to the nature of LISA data, being overtly compatible with estimation of the nonstationary detector noise [28,29], as well as analysis approaches in the presence of data gaps [30,31]. While promising in principle, however, we will see that the main benefit of semicoherent statistics such as S_N (a broadened peak around θ^*) is slightly negated by the retained susceptibility to noise (which is on par with that of X).

Rather than performing multiple realizations of full-scale Markov-chain Monte Carlo simulations, we may gain some insight into the *intrinsic traversability* of various objective functions by simply examining the typical speed of a directed Metropolis–Hastings chain along the connecting line from secondary to injection. This approach is used to conduct a ceteris-paribus comparison of the noisy ($x = h^* + n$) functions shown in the bottom panel of Fig. 3, with minimal and identical assumptions on the sampling algorithm that is used. It is clear that X itself is significantly more difficult to traverse than the other functions, so we shall anneal it by an effective temperature of ρ^* to ensure a fair fight. Other functions have also been proposed in the literature (or are currently under investigation), both generally and specifically for EMRIs [9,32,33]. We do not consider these here, but only note that the same set of criticisms raised in this manuscript are likely to apply (they are either overly *ad hoc*, or susceptible to noise).

In this simple directed-chain analysis, a Markov chain is constrained to lie on the connecting line $\{\theta_i\}$ from θ^S ($i = 1$) to θ^* , with a fixed-step transition $\pm(\theta_{i+1} - \theta_i)$ that is small enough to resolve the fine structure in all

connections (i.e., the variations due to noise). The proposal distribution is

$$P_p(\theta_j|\theta_i) = \begin{cases} 1 - 1/(G+1) & j > i \\ 1/(G+1) & i < j \end{cases} \quad (36)$$

where $G > 1$ is a “diffusion gradient” of sorts, and the acceptance probability for each objective function F is

$$P_A(\theta_j|\theta_i; F) = \min \left\{ 1, \frac{1}{G} \exp(F(\theta_j) - F(\theta_i)) \right\}. \quad (37)$$

Each Markov chain is allowed a maximum length of 10^3 , and 10^3 such chains are produced for the four functions X/ρ^* , S_6 , S_{12} and f . The results are summarized in Fig. 4, where the minimum, median, and maximum locations of the chain (i.e., the indices of θ_i) are shown at each step in the chain. Perhaps unsurprisingly, the median chain speed for each objective function is dictated by the degree of noise in the corresponding connections; the least-traversable function is S_6 , followed by S_{12} (which is visually smoother from the bottom panel of Fig. 3), then X/ρ^* (which retains the secondary peak but is even smoother in the tail region) and finally f (which is virtually constant in the tail region). This trend does not appear to depend on the diffusion gradient, with the speed order preserved at various examined values of G .

The above analysis is not fully conclusive, of course, since the chains are directed along a line toward the global peak for the sake of efficiency. In a higher-dimensional setting, a search with f will take longer to chance upon the peak, whose extent *at the baseline value* is determined by the local length scale of fully coherent signal correlations (the same holds for X and $\ln L$, regardless of any annealing factor). This is not nearly as localized as the *posterior bulk*

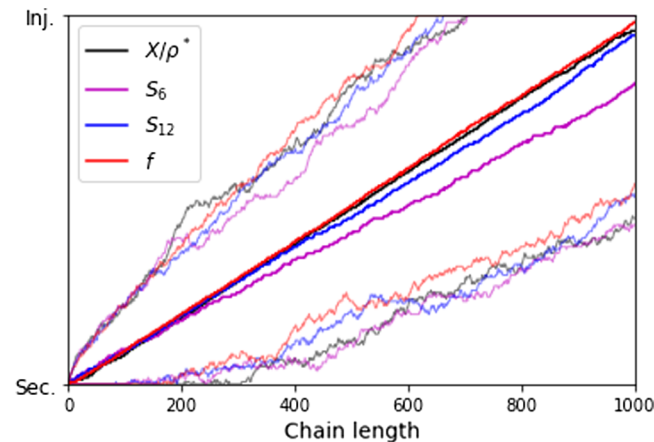


FIG. 4. Minimum, median and maximum locations of 10^3 chains for each objective function (with maximum length 10^3), directed along the connecting line from secondary to injection. The diffusion gradient is $G = 2$, chosen such that approximately half of the chains for f reach the injection parameters.

itself, whose characteristic length scale is inversely proportional to the injection SNR; see Sec. III C. Thus the global peak in f is no more difficult to find than its counterpart in, say, the highest-temperature likelihood of a parallel-tempering algorithm. On the other hand, while the larger “capture region” of the semicoherent statistic will no doubt be beneficial, strong nonlocal variations due to signal correlations are still present in its tail region (this is not evident along the connections shown in Fig. 3, but the baseline is certainly not ≈ -5 across the signal space). Whether or not these larger-scale gradients will be a net help or hindrance to semicoherent search remains to be determined.

It is also possible to construct f -based functions with broadened peaks by conducting less informative comparisons between data and template, but without resorting to the partial maximization or marginalization of the function over some degrees of freedom. The most direct approach is simply to alter the signal space altogether, by truncating the signal templates (and data) to some shorter observation duration with the original reference initial time (defined implicitly by (p_0, e_0, t_0)). An example of f for a truncated duration of three months is shown in Fig. 3 (dashed red connections), where the broadening of the primary peak and the retained suppression of all other variations are both evident. The function f with a range of truncated observation durations might thus be used in a sampling scheme that is akin to parallel tempering, or similar hierarchical-style searches.

We have thus far been considering only the single-signal hypothesis H_1 for identification. In the case of multiple EMRI signals, the difficulty in searching the signal space is amplified for all functions. However, it is not hard to see that the complexity of f will only be weakly affected (unless secondaries from different signals are likely to combine, but an analytical calculation in Chua and Cutler [9] indicates that they are not). Strong peaks in f exist only at the parameters of actual signals; while search chains will tend to be trapped at each location, this does not detract from either parallelized or sequential search, and is even beneficial as it facilitates a direct transition into inference (see Sec. III C).

The efficiency of sampling algorithms when used for optimization relies not only on the global structure of the objective function, but also the intrinsic cost of evaluating the function. While the evaluation of the vector-valued detection statistic \mathbf{X} incurs an overt factor of M in computational complexity over X (or L), this only multiplies the cost of the inner-product operation $\langle \cdot | \cdot \rangle$. The cost of each such operation certainly dominates the total cost on the signal-processing end—but the EMRI template h itself is also notoriously expensive to generate without computational enhancements, and ultimately even irreducibly so due to its sheer length and complexity. Depending on the efficiency of the template model, the function f should thus

be evaluable at little to modest additional computational cost over X or L .

C. As a likelihood function

Finally we turn to the procedure of inference, by examining the utility of $f|H_1$ as an *effective* (log-)likelihood function. Recall that at leading order, the standard log-likelihood (8) in the vicinity of θ^* is simply [34]

$$\ln L = -\frac{1}{2}\Delta^T \mathcal{I} \Delta + \mathcal{O}(|\Delta|^3), \quad (38)$$

where $\Delta := \theta - \theta^*$, and the Fisher information \mathcal{I} is given component-wise by the pullback metric on signal space:

$$\mathcal{I} = \langle \partial_\theta h | \partial_\theta h \rangle |_{\theta=\theta^*}. \quad (39)$$

At a sufficiently high SNR, the full likelihood becomes well represented by Eq. (38), and the resultant posterior appears close to Gaussian in the bulk region. For the sake of concreteness, we will define the posterior-bulk region here as the set of points where $\ln L$ exceeds the 3-sigma value $-9/2$ for the leading-order likelihood (with a locally uninformative prior).

Assuming the components of $\partial_\theta \rho^2 = 2\langle h | \partial_\theta h \rangle$ and $\partial_\theta \beta$ are negligible in the bulk (which is generally valid for all parameters but the luminosity distance), a similar expansion of f about θ^* yields

$$f = -\frac{1}{2}\Delta^T \mathcal{I}' \Delta + \mathcal{O}(|\Delta|^3), \quad (40)$$

where

$$\mathcal{I}' = \beta \rho \left(\mathcal{I} - \sum_m \mathcal{I}_m \right) + \frac{1}{\rho} \mathcal{I}, \quad (41)$$

with the ‘‘Fisher information’’ for each mode given by

$$\mathcal{I}_m = \langle \partial_\theta h_m | \partial_\theta h_m \rangle |_{\theta=\theta^*}. \quad (42)$$

The exponential factor in Eq. (1) contributes to the first term in Eq. (41), while the second term is proportional to the Fisher information and arises from the expansion of X about θ^* (with the additional factor of $1/\rho$ due to the template normalization in X).

Positive-definiteness induces a partial ordering on the set of all symmetric matrices (and thus the subset of positive-definite matrices as well). We now assume that $\mathcal{I}_m < \mathcal{I}_{m+1} < \mathcal{I}$ for all m . Then (very roughly):

$$\mathcal{I} - \sum_m \mathcal{I}_m \approx \mathcal{I} - \mathcal{I}_M \approx (1 - \alpha^2) \mathcal{I}, \quad (43)$$

and we have

$$\mathcal{I}' \approx \frac{2 \ln(\alpha \rho) + 1}{\rho} \mathcal{I}. \quad (44)$$

Since the decomposition (2) for EMRIs is in terms of signal-like modes with the same underlying dependence on the source parameters, these modes vary in qualitatively similar ways to the full signal template, such that Eq. (44) is not a terrible approximation. In fact, the local profile of f about θ^* generally looks consistent with negative excess kurtosis, in which case the approximation is actually conservative. For $\rho > 5$ and all α , Eq. (44) indicates a broadening of f over $\ln L$, since we have

$$\sqrt{\frac{\rho}{2 \ln(\alpha \rho) + 1}} > 1, \quad (45)$$

where the left-hand side is the associated length-scale factor. This quantity is plotted as a function of α in Fig. 5, for various values of ρ^* [it does not depend on M due to the crude approximation in Eq. (43)].

For our representative signal injection, $\alpha \approx 0.9$ and the estimated length-scale factor is ≈ 1.7 —this is a slight underestimate of the actual broadening of f over $\ln L$, which is by a factor of between two and three (see Fig. 6). Also included in Fig. 6 are the profiles of f' from Eq. (21) (which, recall, is slightly more flat-topped and locally closer to X), as well as X itself (where the broadening over $\ln L$ is by a factor of $(\rho^*)^{1/2} \approx 4.5$, as expected).

With the general broadening of f over $\ln L$, the extent of the recovered distribution from the sampling of f provides useful prior localization for regular posterior estimation with L . More promisingly, note from Fig. 5 that the broadening is also modest (no more than a factor of five at an SNR of 20). This constrained broadening of f over $\ln L$ implies that the search chains might (upon a suitable level of convergence) be used *directly as posterior samples*, with the unnormalized importance weights

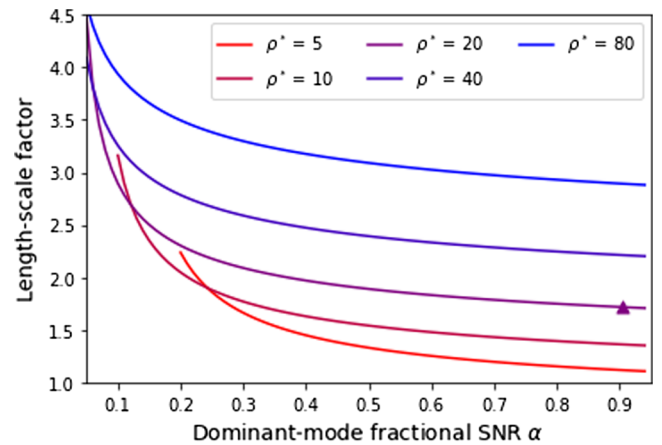


FIG. 5. Dependence of the estimated length-scale factor (45) on α , for various ρ^* . The triangle corresponds to the representative injection considered in this work, for which $\alpha \approx 0.9$.




FIG. 6. Extended connections of $\ln L$, X , f and f' between injection and secondary parameters, with noise and zoomed in on (a one-dimensional slice of) the posterior-bulk region.

$$w := L\pi \exp(-f), \quad (46)$$

where π is the desired Bayesian prior.

IV. CONCLUSION

In this manuscript, I propose a potentially useful one-stop function $f(\theta)$ for various tasks in GW data analysis; it is conditioned on the detector data x , and defined for some signal-template model $h(\theta)$ with a general mode decomposition. The basic mechanics of f rely critically on a set of assumptions about the decomposition—essentially that signal templates may readily be written as the sum of $M > 1$ uncorrelated modes, each with substantial power. This work is primarily motivated by the deep-rooted problem of *strong nonlocal parameter degeneracy* in the space of EMRI signals, and so recent results by Chua and Cutler [9] on the nature of this phenomenon are used to define a model-specified calibration of the function for EMRI data analysis. I then build a case for the utility of f by examining its properties as a statistic for detection, as an objective function for identification, and as an effective likelihood function for inference.

The main difficulty in EMRI data analysis is the procedure of source identification, which is essentially a large-scale optimization problem. Traditionally defined objective functions over the model space, such as the standard detection statistic or standard likelihood function in GW data analysis, suffer from the presence of numerous and highly pronounced secondary peaks that hinder identification. The function proposed here is based on the principle of *deemphasizing* these peaks in some way, which is shared by well-known strategies such as semicoherent filtering or annealing-type sampling algorithms. Semicoherent filtering also falls under another class of strategies (including, e.g., F -statistic searches [32]) that use partial maximization or marginalization to conduct less informative comparisons between data and template. This is useful

for EMRI identification in another way, since it broadens all peaks in the function—including the primary peak containing the global maximum, which is highly localized relative to the model space.

I posit a different tenet here: that a function with virtually no gradients is easier to search than one with uncontrolled variations at both large and small length scales. This is realized in the proposed function by using *exponential suppression* to de-emphasize secondaries, rather than congealing them (as in the case of semicoherent filtering) or simply rescaling the entire function (annealing). As a by-product, any less-pronounced variations in the function due to signal correlations or detector noise are suppressed to a near-constant baseline. I argue in Sec. III B that the elimination of such variations might play a larger role in the efficacy of search algorithms than any other individual factor, although the jury is still out on whether exponential suppression is generally superior to partial maximization or vice versa. To provide a useful (if somewhat biased) analogy, the latter is akin to searching for a bigger needle in a haystack (really a field of haystacks), while the former only involves finding a smaller needle in an open field.

The “one-stop” aspect of the proposed function also connects identification to detection and inference more seamlessly, at least at the conceptual level. In practice, of course, one must still assess the statistical significance of signal candidates through traditional methods, and will likely take the effort to repeat inference on identified signals directly with the standard likelihood—but neither is intrinsically challenging to do for EMRIs. As briefly discussed in Sec. III B, the function can easily be used in hierarchical approaches to replicate the advantages of partial maximization. It is also viable in the broader context of the LISA global fit: search and inference for individual (resolvable) EMRI signals will almost certainly be performed independently using the catalog residuals, as they do not strongly impact noise estimates, or inference on signals of a different source type [35] (or even on one another [9]). That being said, some thought might be required to adapt the function for the presence of data gaps. Beyond the setting of EMRI data analysis, the general principle of exponential suppression should also, at minimum, provide an interesting alternative to the existing paradigms in coherent GW searches.

ACKNOWLEDGMENTS

I am grateful to Curt Cutler for seminal discussions, and for our closely related collaboration on [9]. Lorenzo Speri suggested the effective broadening of the proposed function with truncated observation durations, for which I am deeply appreciative. Thanks also go out to Katerina Chatziioannou, Neil Cornish, Michael Katz, Christopher Moore and Michele Vallisneri for their insightful input on the manuscript. This work was supported by the NASA LISA Preparatory Science Grant No. 20-LPS20-0005.

- [1] K. Danzmann *et al.*, [arXiv:1702.00786](#).
- [2] P. Amaro-Seoane, *Living Rev. Relativity* **21**, 4 (2018).
- [3] 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).
- [4] 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](#).
- [5] L. Barack and A. Pound, *Rep. Prog. Phys.* **82**, 016904 (2019).
- [6] A. Pound and B. Wardell, Black hole perturbation theory and gravitational self-force, in *Handbook of Gravitational Wave Astronomy*, edited by C. Bambi, S. Katsanevas, and K. D. Kokkotas (Springer, Singapore, 2021).
- [7] A. J. K. Chua, M. L. Katz, N. Warburton, and S. A. Hughes, *Phys. Rev. Lett.* **126**, 051102 (2021).
- [8] M. L. Katz, A. J. K. Chua, L. Sperr, N. Warburton, and S. A. Hughes, *Phys. Rev. D* **104**, 064047 (2021).
- [9] A. J. K. Chua and C. J. Cutler, [arXiv:2109.14254](#).
- [10] J. R. Gair, I. Mandel, and L. Wen, *Classical Quantum Gravity* **25**, 184031 (2008).
- [11] S. Babak, J. R. Gair, and E. K. Porter, *Classical Quantum Gravity* **26**, 135004 (2009).
- [12] S. Babak *et al.* (Mock LISA Data Challenge Task Force Collaboration), *Classical Quantum Gravity* **27**, 084009 (2010).
- [13] N. J. Cornish, *Classical Quantum Gravity* **28**, 094016 (2011).
- [14] Y. Wang, Y. Shang, S. Babak, Y. Shang, and S. Babak, *Phys. Rev. D* **86**, 104050 (2012).
- [15] C. Cutler and E. E. Flanagan, *Phys. Rev. D* **49**, 2658 (1994).
- [16] B. Allen, *Phys. Rev. D* **71**, 062001 (2005).
- [17] S. A. Hughes, N. Warburton, G. Khanna, A. J. K. Chua, and M. L. Katz, *Phys. Rev. D* **103**, 104014 (2021).
- [18] J. Abadie *et al.* (LIGO Scientific, VIRGO Collaborations), *Phys. Rev. D* **85**, 082002 (2012).
- [19] A. J. K. Chua and J. R. Gair, *Classical Quantum Gravity* **32**, 232002 (2015).
- [20] A. J. K. Chua, C. J. Moore, and J. R. Gair, *Phys. Rev. D* **96**, 044005 (2017).
- [21] C. Cutler, *Phys. Rev. D* **57**, 7089 (1998).
- [22] L. Barack and C. Cutler, *Phys. Rev. D* **69**, 082005 (2004).
- [23] C. J. Moore, D. Gerosa, and A. Klein, *Mon. Not. R. Astron. Soc.* **488**, L94 (2019).
- [24] Y.-A. Ma, Y. Chen, C. Jin, N. Flammarrion, and M. I. Jordan, *Proc. Natl. Acad. Sci. U.S.A.* **116**, 20881 (2019).
- [25] D. G. Blair, *The Detection of Gravitational Waves* (Cambridge University Press, Cambridge, 1991).
- [26] T. Damour, B. R. Iyer, and B. S. Sathyaprakash, *Phys. Rev. D* **57**, 885 (1998).
- [27] J. R. Gair, L. Barack, T. Creighton, C. Cutler, S. L. Larson, E. S. Phinney, and M. Vallisneri, *Classical Quantum Gravity* **21**, S1595 (2004).
- [28] M. C. Edwards, P. Maturana-Russel, R. Meyer, J. Gair, N. Korsakova, and N. Christensen, *Phys. Rev. D* **102**, 084062 (2020).
- [29] N. J. Cornish, *Phys. Rev. D* **102**, 124038 (2020).
- [30] Q. Baghi, I. Thorpe, J. Slutsky, J. Baker, T. Dal Canton, N. Korsakova, and N. Karnesis, *Phys. Rev. D* **100**, 022003 (2019).
- [31] K. Dey, N. Karnesis, A. Toubiana, E. Barausse, N. Korsakova, Q. Baghi, and S. Basak, *Phys. Rev. D* **104**, 044035 (2021).
- [32] P. Jaranowski, A. Krolak, and B. F. Schutz, *Phys. Rev. D* **58**, 063001 (1998).
- [33] Y. Wang, [arXiv:1402.6211](#).
- [34] M. Vallisneri, *Phys. Rev. D* **77**, 042001 (2008).
- [35] E. Racine and C. Cutler, *Phys. Rev. D* **76**, 124033 (2007).