

This is the accepted manuscript made available via CHORUS. The article has been published as:

Beyond the Fisher-Matrix Formalism: Exact Sampling Distributions of the Maximum-Likelihood Estimator in Gravitational-Wave Parameter Estimation

Michele Vallisneri

Phys. Rev. Lett. **107**, 191104 — Published 4 November 2011

DOI: [10.1103/PhysRevLett.107.191104](https://doi.org/10.1103/PhysRevLett.107.191104)

Beyond Fisher: exact sampling distributions of the maximum-likelihood estimator in gravitational-wave parameter estimation

Michele Vallisneri¹

¹*Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109*

In gravitational-wave astronomy, data analysts often wish to characterize the expected parameter-estimation accuracy of future observations. The semianalytical Fisher-matrix formalism provides lower bounds on the spread of the maximum-likelihood estimator across noise realizations; it also provides the leading-order width of the posterior probability density. However, the Fisher matrix is limited to the high-SNR limit, which is often not realized in practice. By contrast, Bayesian-inference Monte Carlos characterize the full posterior as a function of noise for any SNR, but they are usually too computationally costly to repeat for a representative sample of noise realizations. In this article, I describe a semianalytical technique to exactly map the sampling distribution of the maximum-likelihood estimator across noise realizations at any SNR, generating an intermediate amount of information with significantly less computation. This technique can be applied to any estimation problem characterized by additive Gaussian noise.

PACS numbers: 04.30.Db, 04.25.Nx, 04.80.Nn, 95.55.Ym

The first direct detections of gravitational-wave (GW) signals are likely to be achieved in the second half of this decade with second-generation ground-based interferometric detectors such as Advanced LIGO and Virgo [1], which are sensitive to high-frequency GWs (10–1000 Hz). Third-generation instruments such as the Einstein Telescope [2] promise much greater reach and yield, while space-based observatories similar to LISA [3] will be sensitive to low-frequency GWs (10^{-5} – 10^{-1} Hz) in a band populated by thousands of detectable sources in the Galaxy and far beyond. A survey of the scientific literature on GW data analysis in this pre-detection era reveals a few dominant *genres*: articles on data-analysis methods and implementations [e.g., 4]; non-detection and upper-limit analyses of actual collected data [e.g., 5]; and “prospects” papers that examine combinations of GW sources and detectors to characterize the expected rates of detection and accuracies of source-parameter estimation [e.g., 6].

Papers in this last class, especially when they are concerned with Bayesian inference for *modeled* GW signals, follow a well-rehearsed structure: a) derive the GW signal h as a function of the source parameters θ ; b) fix fiducial true values θ_{tr} for them; c) simulate a realization n of detector noise, usually assumed as additive and Gaussian; d) derive the noise-dependent probability distribution $p(\theta;n)$ for the source parameters given the data $s = h_{\text{tr}} + n$; e) characterize the error and uncertainty of $p(\theta;n)$; f) finally, repeat steps c, d, and e for a sufficiently broad sample of noise realizations, since each of these will result in different parameter estimates and uncertainties. Optimally, repeat for different fiducial θ_{tr} .

Step d [calculating $p(\theta;n)$] is usually very computationally expensive: for more than two or three source parameters, in Bayesian inference it must be performed with a stochastic technique such as Markov Chain Monte Carlo [7], which may involve $\sim 10^6$ evaluations of the likelihood for different θ s. Each evaluation requires obtaining the GW signal $h(\theta)$ and performing an FFT (say, for 10^5 points) to compute the likelihood of the noise residual [see the discussion around Eq. (1)]. Step f is therefore a Monte Carlo of Monte Carlos (perhaps

again 10^6 of them), with a total computational cost, for a single θ_{tr} , of $\sim 10^{18}$ floating-point operations, more than can usually be procured easily. More emphatically, I call this scheme the Holy Grail of parameter-estimation prospects: even the papers that try hardest [e.g., 8], perform meta-Monte Carlos of ~ 100 combinations of noise realizations and fiducial sources.

Most “prospects” papers, instead, take advantage of the fact that steps d to f can be short-circuited when the signal-to-noise ratio (SNR) of detection is sufficiently high that $p(\theta;n)$ collapses to a normal distribution centered around the maximum-likelihood parameters $\theta_{\text{ml}}(n)$, with covariance given by the inverse Fisher matrix, which is *not* a function of n . At high SNR, the distribution of θ_{ml} for different noise realizations is itself normal and centered at θ_{tr} , with covariance again given by the inverse Fisher matrix. The computational cost of this approximation (for a single θ_{tr} and 10^5 -point likelihoods) is $\sim d^2 \times 10^6$ floating-point operations, where d is the number of source parameters, affording virtually instantaneous results.

Unfortunately, as I discuss at length in Ref. [9] and as highlighted elsewhere [6, 10], for many practical parameter-estimation problems SNRs will not be sufficiently high to justify the approximation. The crux of the problem is that the Fisher matrix, built from the partial derivatives $\partial h / \partial \theta_i$ of the signal, can only represent h correctly if h is *linear* in all the θ_i across ranges comparable to the expected parameter errors. These decrease as SNR grows, making the condition less stringent; in Ref. [9] I provide a criterion to determine when SNR is high enough.

Thus, in general a Monte Carlo is needed for step d; furthermore, the meta-Monte Carlo of step f is also necessary, because the particular noise realization *does* affect parameter uncertainties, as shown in Fig. 2 for the toy model discussed later in this paper. It is true [11] that for additive Gaussian noise the average of $p(\theta;n)$ across all ns equals $p(\theta;n=0)$; but this averaged likelihood describes an unrealistic limit (infinite observations of the same source with finite total SNR) that is not representative of average or typical errors, except trivially in the high-SNR limit. Thus, the shape and width of generic

$p(\theta; n)$ cannot generally be obtained from $p(\theta; n = 0)$, as suggested in Ref. [12]. In Ref. [9] I derive an expansion of $p(\theta; n)$ in powers of $1/\text{SNR}$, where n -dependent terms are seen at next-to-leading order for *both* the centering and shape of the posterior.

In this paper I describe and test a technique to perform as less thorough, but more affordable survey than the Monte Carlo of Monte Carlos: mapping the sampling distribution of the maximum-likelihood (henceforth, ML) estimator θ_{ml} [the maximum of each $p(\theta; n)$] over all noise realizations. From a classical-statistics (a.k.a. “frequentist”) viewpoint, the spread of this distribution is formally the uncertainty of the ML point estimator. From a Bayesian viewpoint [13], if priors are unimportant, the spread of this distribution describes one major component of the expected error (the other is the intrinsic spread of the posterior for each n).

Mapping the sampling distribution of the ML estimator.

We consider a large set of experiments in which we observe the signal $h(\theta_{\text{tr}})$, where θ_{tr}^i is the d -dimensional vector of true parameter values and h is an N -dimensional vector (e.g., a time series) with $N \gg d$. In each experiment, the detector output is $s = n + h(\theta_{\text{tr}})$, where n is additive Gaussian noise distributed with $p(n) = \mathcal{N} \exp[-(n, n)/2]$. Here (\cdot, \cdot) is the standard signal inner product, given in one convention [6] by $(s, t) = 4 \text{Re} \int_0^\infty \tilde{s}^*(f) \tilde{t}(f) / S(f) df$, with $\tilde{\cdot}$ denoting the Fourier transform and $*$ the complex conjugate. In this paper, without loss of generality, we treat (\cdot, \cdot) as the inner product of an abstract linear space, and let $|s|^2 \equiv (s, s)$.

For a given θ_{tr} and noise realization n_{tr} , the ML estimator $\theta_{\text{ml}}(n_{\text{tr}}, \theta_{\text{tr}})$ maximizes the probability of the residual noise n obtained by subtracting the postulated signal $h(\theta)$ from the data $s = h(\theta_{\text{tr}}) + n_{\text{tr}}$; that is, it maximizes

$$p(n = s - h(\theta)) \propto e^{-|h(\theta_{\text{tr}}) + n_{\text{tr}} - h(\theta)|^2/2}. \quad (1)$$

Thus, θ_{ml} must satisfy the vector equation

$$\text{ML}_i(\theta; n, \theta_{\text{tr}}) \equiv (\partial_i h(\theta), h(\theta_{\text{tr}}) + n - h(\theta)) = 0, \quad (2)$$

where the ML_i are the partial derivatives of $-2 \log p$.

Our purpose is to map the distribution of θ_{ml} across all noise realizations. We can do this by enumerating all possible n s (weighted by $p(n)$), figuring out the θ_{ml} corresponding to each, and accumulating the resulting distribution of θ_{ml} . Formally,

$$p(\theta_{\text{ml}} = \theta | \theta_{\text{tr}}) = \int \delta(\theta_{\text{ml}}(n, \theta_{\text{tr}}) - \theta) p(n) dn. \quad (3)$$

Unfortunately, because $h(\theta)$ is generally a complicated function of the θ^i , it is difficult to solve $\text{ML}_i = 0$ for θ_{ml} given n , so we can only integrate Eq. (3) using a Monte Carlo approach where we generate full, high-dimensional realizations of n (e.g., as time series), search parameter space for θ_{ml} by repeatedly evaluating Eq. (1), and iterate for many different n .

The main result of this paper is that there *is* a more effective way to map $p(\theta_{\text{ml}})$: we enumerate the θ_{ml} , and compute the total probability weight of the n that are compatible with each. (Put slightly differently, for each θ_{ml} we

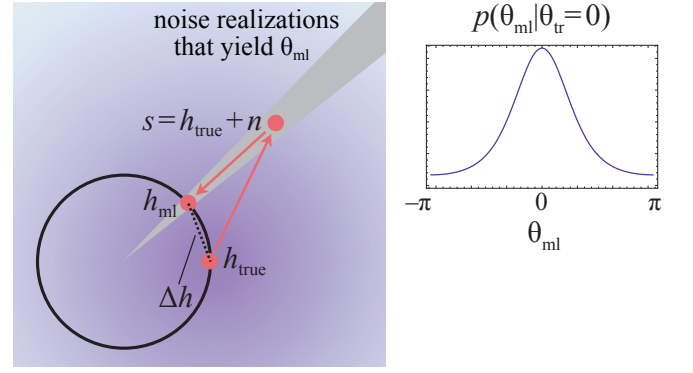


FIG. 1. To compute $p(\theta_{\text{ml}})$, we integrate $p(n)$ over all noise realizations that satisfy $\text{ML}_i = 0$ —see discussion below Eq. (4). Here $h(\theta) = (\cos \theta, \sin \theta)$, $p(n_x, n_y) = (2\pi)^{-1} \exp[-(n_x^2 + n_y^2)/2]$, and $\theta_{\text{ml}} = \arctan(s_y/s_x)$. The resulting $p(\theta_{\text{ml}} | \theta_{\text{tr}} = 0)$ is shown top right.

count how many experiments would yield it as the maximum of likelihood.) Because the ML_i are *linear* functions of the n , this weight is the probability mass of the $(N - d)$ -dimensional subspace of the noise vectors that solve $\text{ML}_i = 0$ for all i . Using the δ -of-function relation for the vector ML_k , $\delta(\text{ML}_k(\theta; n, \theta_{\text{tr}})) = \delta(\theta_{\text{ml}}(n, \theta_{\text{tr}}) - \theta) / |\partial \text{ML}_i(\theta) / \partial \theta_j|$, we can then rewrite $p(\theta_{\text{ml}} = \theta | \theta_{\text{tr}})$ as

$$\mathcal{N} \int \Pi_k \delta(\text{ML}_k(\theta; n, \theta_{\text{tr}})) \times |\partial \text{ML}_i / \partial \theta_j| e^{-(n, n)/2} dn. \quad (4)$$

Figure 1 exemplifies the process of integrating over the n compatible with a chosen θ_{ml} . In this low-dimensional model [14], detector data is described by a point in the plane, and the signal family $h(\theta)$ by the unit circle. For each n , the true signal h_{true} is displaced to a different $s = h_{\text{true}} + n$ with probability $p(n)$ indicated by the shading. Projecting back to $h(\theta)$ identifies the ML waveform h_{ml} and parameter θ_{ml} . All noise realizations that produce an s within the gray sector project to the same h_{ml} ; thus, integrating $p(n)$ over the sector yields $p(\theta_{\text{ml}})$.

Equation (4) is especially powerful because, contrary to appearance, it does not require integration over the full N dimensions of the noise, but only over $\sim d^2$ coordinate directions corresponding to the ML_k and $\partial \text{ML}_i / \partial \theta_j$. Since both these sets of functions are linear in n , they can be seen as jointly normal random variables that are fully characterized by their inner products; all the other noise degrees of freedom have no effect on the integral other than its normalization. (We spell this out in detail in the next section.) With Eq. (4) in hand, we can then sample $p(\theta_{\text{ml}})$ directly (for low d), or by Markov Chain Monte Carlo techniques.

Evaluating the master integral. Equation (4) can be evaluated elegantly as the expectation value of a function (the determinant) of correlated random variables. For clarity, we follow a more pedestrian approach: we begin by transforming the N -dimensional integral over the n to new coordinates where the first few basis vectors span the random variables of interest. Namely, we write n in terms of a new basis where the first d

vectors are obtained by orthonormalizing the $\partial_i h \equiv h_i$,

$$\hat{n}_1 \propto h_1, \quad \hat{n}_i \propto \left(1 - \sum_{j=1}^{i-1} \hat{n}_j \otimes \hat{n}_j\right) \cdot h_i; \quad (5)$$

furthermore, we obtain the $(d+1)$ -th to $d(d+3)/2$ -th basis vectors by orthonormalizing the $\partial_{ij} h \equiv h_{i,j}$ with $i \geq j$; the remaining $N - d(d+3)/2$ noise combinations complete the basis.

We can now rewrite all the variables that appear in Eq. (4) in terms of the new basis:

$$n \equiv N^k \hat{n}_k, \quad h_i \equiv C_i^k \hat{n}_k, \quad h_{i,j} \equiv C_{ij}^k \hat{n}_k; \quad (6)$$

here the N^k , C_i^k , and C_{ij}^k are the coefficients that expand the n , h_i , and $h_{i,j}$ in terms of the \hat{n}_k . We use Einstein summations, and treat (ij) both as a double index and as a single index flattened to the range $(d+1), \dots, d(d+3)/2$. By virtue of orthonormalization, $C_i^k = 0$ for $k > i$ and $C_{ij}^k = 0$ for $k > (ij)$, with $C_i^i > 0$ and $C_{ij}^{(ij)} > 0$. Furthermore,

$$\begin{aligned} \text{ML}_i &\equiv C_i^k N_k - C_i^k (\Delta h, \hat{n}_k), \\ \text{ML}_{i,j} &\equiv C_{ij}^m N_m - C_{ij}^m (\Delta h, \hat{n}_m) - C_i^k C_{jk}, \end{aligned} \quad (7)$$

with $\Delta h \equiv h(\theta) - h(\theta_{\text{tr}})$, and with the sums over k limited to $k \leq i$ and the sums over m limited to $m \leq (ij)$.

To perform the integral in Eq. (4), we use the first δ to fix $N_1 = (\Delta h, \hat{n}_1)$, yielding a δ -normalization factor of $1/C_1^1$; we use the second δ to fix $N_2 = (\Delta h, \hat{n}_2)$ (after the terms proportional to N_1 in the δ cancel out), yielding a factor of $1/C_2^2$; and so on. Next, we perform the trivial integral over the “signal-orthogonal” noise degrees of freedom that correspond to the \hat{n}_k with $k > d(d+3)/2$, leaving

$$p(\theta_{\text{ml}} = \theta | \theta_{\text{tr}}) = \frac{e^{-\sum_{k=1}^d (\Delta h, \hat{n}_k)^2 / 2}}{(2\pi)^{d(d+3)/4} \prod_{k=1}^d C_k^k} \times \int \left| C_{ij}^m N_m - C_{ij}^m (\Delta h, \hat{n}_m) - C_i^k C_{jk} \right| e^{-N^m N_m / 2} dN_m; \quad (8)$$

here the sums over m span $[15] m = (d+1), \dots, d(d+3)/2$, and the sum over k spans $k = 1, \dots, d$.

For a given θ , the main computational cost of evaluating Eq. (8) resides in the orthonormalization and in the computation of the $(\Delta h, \hat{n}_m)$, which together require $\sim d^4/8$ inner products (and therefore N -dimensional signal FFTs); by contrast, the $d(d+1)/2$ -dimensional Gaussian integral can be evaluated much more cheaply (e.g., by a 10,000-point Monte Carlo over N_m drawn from a normal distribution), since the integrand is a function of small matrices and not long FFTs.

Fisher-matrix limit. The well-known high-SNR, Fisher-matrix limit, in which the waveform can be approximated as a linear function of the parameters (and in which the θ_{ml}^i have a simple jointly-normal distribution), follows easily by specializing Eq. (8). Without loss of generality, we set $h(\theta) = \theta^i h_i$; it follows that $h_{i,j} = C_{ij}^k = 0$. The integration

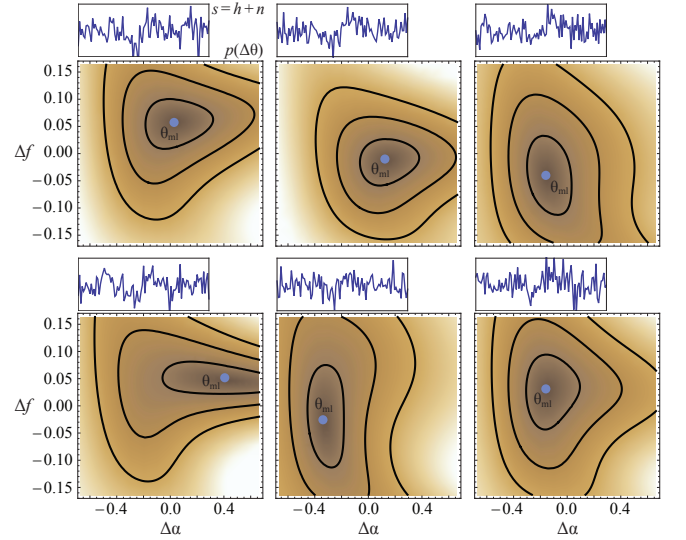


FIG. 2. Likelihood maps obtained for the toy model (11) with different n s and with the same $(\alpha_{\text{tr}}, f_{\text{tr}})$, located at the origin of the plots. Contours are plotted at levels that would correspond to 1- σ , 2- σ , and 3- σ ellipses for a normal distribution. The signal $s = h + n$ is shown above each map.

over the N_m is trivial, and yields $(2\pi)^{d(d-1)/4} |C_i^k C_{jk}|$, where $C_i^k C_{jk} = (h_i, h_j) \equiv F_{ij}$ is the infamous Fisher matrix. Furthermore, $\Delta h = h_i \Delta \theta^i$, where $\Delta \theta^i \equiv \theta_{\text{ml}}^i - \theta_{\text{tr}}^i$ is the error of the ML estimator, so the first exponential of Eq. (8) can be rewritten as $\exp(-\Delta \theta^i F_{ij} \Delta \theta^j / 2)$, since the \hat{n}_k , for $k = 1, \dots, d$, span a complete basis for the h_i . Last, because of the structure of the C_i^k , $\prod_{k=1}^d C_k^k = \sqrt{|C_i^k C_{jk}|} = \sqrt{|F_{ij}|}$. Taking everything together, we reproduce the Fisher-matrix result for the distribution of θ_{ml} [16],

$$p(\theta_{\text{ml}} | \theta_{\text{tr}}) = \frac{e^{-\Delta \theta^i F_{ij} \Delta \theta^j / 2}}{\sqrt{(2\pi)^d |F_{ij}^{-1}|}}. \quad (9)$$

The general result [Eq. (8)] can be restated in terms of F_{ij} , in a form more suitable to computation:

$$p(\theta_{\text{ml}} = \theta | \theta_{\text{tr}}) = \frac{e^{-(\Delta h, h_i) (F^{-1})^{ij} (\Delta h, h_j) / 2}}{\sqrt{(2\pi)^d |F_{ij}|} \sqrt{(2\pi)^{d(d-1)/2} |D_{\mu\nu}|}} \times \int |F_{ij} + (\Delta h, h_{ij}) - M_{(ij)}| e^{-M_\mu (D^{-1})^{\mu\nu} M_\nu / 2} dM_\mu; \quad (10)$$

here $D_{\mu\nu} \equiv D_{(ij)(kl)}$ is the $d(d-1)/2$ -dimensional square matrix given by the products $(h'_{i,j}, h'_{k,l})$ with $j \leq i, l \leq k$: the primes denote projection orthogonal to h_k (i.e., $h'_{i,j} = \sum_{d+1}^{d(d+3)/2} C_{ij}^k \hat{n}_k$); and $M_{(ij)}$ is the matrix obtained from the $d(d-1)/2$ -dimensional vector M_μ of integration variables by remapping indices.

Toy model. To exemplify the use of Eqs. (4) to map $p(\theta_{\text{ml}})$ for low-SNR parameter estimation, let us consider a family of sine-Gaussian signals given by

$$h(t; A, \alpha, f) = A e^{-t^2/2\alpha^2} \sin(2\pi f t). \quad (11)$$

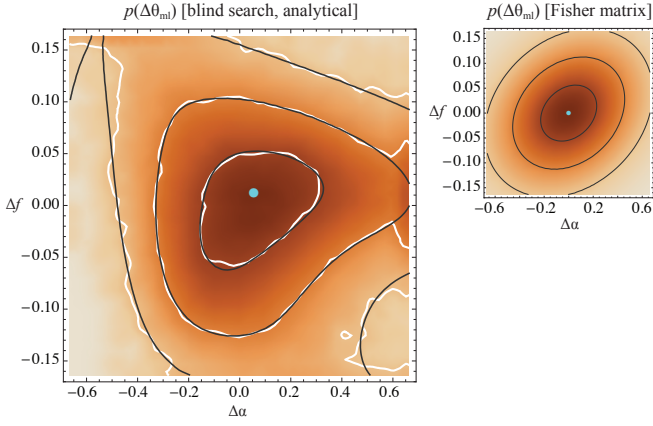


FIG. 3. Exact (left) and Fisher-matrix (right) sampling distribution of θ_{ml} . Contours are plotted as in Fig. 2, with the noisy white curves and shading derived from the numerical maxima of 100,000 likelihood maps for different ns , and the dark curves from Eq. (4). The dot marks the most probable θ_{ml} , slightly displaced from θ_{tr} .

We consider the problem of jointly estimating α and f , but not A , which we fix to yield $\text{SNR} = 5$. For our example, we select $\alpha_{\text{tr}} = 1$ and $f_{\text{tr}} = 0.25$. As shown in Fig. 2, at this low SNR different noise realizations yield strikingly different likelihood maps for the same true signal—all quite different from the ellipsoidal Fisher prediction! In each map, the pale blue dot marks the location of θ_{ml} —indeed, the distribution of blue dots is just the desired $p(\theta_{\text{ml}})$.

In Fig. 3, I show maps of $p(\theta_{\text{ml}})$ obtained using both the brute-force approach [Eq. (3)] and the new method [Eq. (4)]. For the former (the noisy white curves and shading), I produced 100,000 likelihood maps $p(\theta; n)$ with different ns drawn from $p(n)$, and 2D-histogrammed the resulting θ_{ml} . For the latter (dark curves), I used Eq. (4), as implemented by Eq. (8). As expected, the maps agree, but the new method is considerably faster. For comparison, the top-right plot shows the Fisher-matrix prediction (9).

Conclusions. I have described a novel approach to create exact maps, for any SNR, of the distribution of the ML estimator for the source parameters of a signal embedded in additive Gaussian noise. This distribution would be obtained in a large set of observations of the same true signal with different noise realizations, each appearing with probability $p(n)$. Given a single observation, such a map embodies the frequentist notion of uncertainty for the ML estimator. From a Bayesian viewpoint, if priors are unimportant, the map characterizes the distribution of possible maxima of posterior probabilities.

In comparison to the computational cost of the “Holy Grail” Monte Carlo of Monte Carlos (10^{18} operations), we estimate the cost of Eq. (4) as follows: 10^6 Monte Carlo samples of candidate θ_{ml} s, times $d^4/8$ inner products, times 10^6 floating-point operations for each inner-product FFT (again assuming $N = 10^5$); thus, even for $d = 10$, this scheme involves $\sim 10^{15}$ operations—a thousand times cheaper.

These maps can be used directly, in both frequentist and Bayesian frameworks, to study parameter-estimation

prospects, but also to perform stringent tests of Fisher-matrix predictions at low SNRs, and to provide proposal distributions for Monte Carlo searches of unknown sources. An interesting feature in this regard is that $\text{ML}_i = 0$, a local condition, does not distinguish between primary and secondary maxima of the likelihood, and it will include the latter (if they are sufficiently probable) in the maps; thus Eq. (4) could be exploited to enable jumps between separated peaks in complex likelihoods, which are generally very difficult to locate. While this result was derived in the context and with the motivation of GW science, it is applicable to statistical inference for any problem where noise can be regarded as Gaussian and additive, such as several that arise in high-energy physics and observational cosmology.

Acknowledgments. I am grateful to G. Cicuta, N. Cornish, C. Cutler, F. Feroz, M. Hobson, J. Jewell, I. Mandel, S. Nissanke, E. Onofri, R. O’Shaughnessy, and T. Prince, as well as two anonymous referees, for useful suggestions and for reviewing this manuscript. This work was supported by the RTD program at the Jet Propulsion Laboratory, California Institute of Technology, where it was performed under contract with the National Aeronautics and Space Administration. Copyright 2011 California Institute of Technology. Government sponsorship acknowledged.

-
- [1] G. M. Harry and the LIGO Scientific Collaboration, *Class. Quant. Grav.* **27**, 084006 (2010); Virgo Collaboration, “Advanced Virgo baseline design,” tech. report VIR-0027A-09, tds.ego-gw.it/ql/?c=6589 (2009).
- [2] Special issue, *Gen. Rel. Grav.* **43**, 361 (2011).
- [3] T. A. Prince *et al.*, “LISA: Probing the Universe with Gravitational Waves,” list.caltech.edu/mission_documents (2009).
- [4] B. Allen *et al.*, [arXiv:gr-qc/0509116](https://arxiv.org/abs/gr-qc/0509116) (2005).
- [5] B. Abbott *et al.*, *Astrophys. J. Lett.* **683**, L45 (2008).
- [6] C. Cutler and É. E. Flanagan, *Phys. Rev. D* **49**, 2658 (1994).
- [7] J. S. Liu, *Monte Carlo strategies in scientific computing* (Springer, New York, 2001).
- [8] S. Nissanke *et al.*, *Astrophys. J.* **739**, 99 (2011).
- [9] M. Vallisneri, *Phys. Rev. D* **77**, 042001 (2008).
- [10] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar, *Phys. Rev. D* **53**, 3033 (1996); R. Balasubramanian and S. V. Dhurandhar, *ibid.* **57**, 3408 (1998); D. Nicholson and A. Vecchio, *ibid.* **57**, 4588 (1998); S. Vitale and M. Zanolin, *ibid.* **82**, 124065 (2010).
- [11] S. Nissanke *et al.*, *Astrophys. J.* **725**, 496 (2010).
- [12] N. J. Cornish, [arXiv:1007.4820 \[gr-qc\]](https://arxiv.org/abs/1007.4820) (2010).
- [13] E. T. Jaynes, *Probability Theory: The Logic of Science* (Cambridge University Press, 2003).
- [14] Suggested by C. Cutler in a personal communication (2010).
- [15] Inside the determinant, the first d terms of the form $C_{ij}^m[N_m - (\Delta h, \hat{n}_m)]$ vanish because of the δ s applied previously.
- [16] Formally, we see that Eq. (8) tends to Eq. (9) because the terms other than F_{ij} in the determinant are suppressed by a factor $1/\text{SNR}$. This follows from C_i^k and $C_{ij}^m \sim O(\text{SNR})$, N_m and $\hat{n}_m \sim O(1)$, and $\Delta h \sim O(\text{SNR}) \times O(\sqrt{F^{-1}}) \sim O(1)$.