

Targeted search for continuous gravitational waves: Bayesian versus maximum-likelihood statistics

Reinhard Prix¹, Badri Krishnan²

¹ Albert-Einstein-Institut, Callinstr. 38, 30167 Hannover, Germany

² Albert-Einstein-Institut, Am Mühlenberg 1, 14476 Potsdam, Germany

E-mail: Reinhard.Prix@aei.mpg.de

LIGO-P0900066-v2 Fri May 22 17:10:14 2009 +0200
commitID: 676b4bee4decac8f9a2f2d399c3d7f79c468ad9b

Abstract. We investigate the Bayesian framework for detection of continuous gravitational waves (GWs) in the context of targeted searches, where the phase evolution of the GW signal is assumed to be known, while the four amplitude parameters are unknown. We show that the orthodox maximum-likelihood statistic (known as \mathcal{F} -statistic) can be rediscovered as a Bayes factor with an unphysical prior in amplitude parameter space. We introduce an alternative detection statistic (“ \mathcal{B} -statistic”) using the Bayes factor with a more natural amplitude prior, namely an isotropic probability distribution for the *orientation* of GW sources. Monte-Carlo simulations of targeted searches show that the resulting Bayesian \mathcal{B} -statistic is more powerful in the Neyman-Pearson sense (i.e. has a higher expected detection probability at equal false-alarm probability) than the frequentist \mathcal{F} -statistic.

1. Introduction

Searches for gravitational waves (GWs) often consist of testing the data for the presence of signals from a *known family* of waveforms, parametrized by (generally unknown) signal parameters. Here we consider the class of coherent GW signals of constant amplitude and polarization, which include “continuous GWs”, e.g. from non-axisymmetric spinning neutron stars (see [1] for a review), stellar-mass binary systems in the LISA frequency band (e.g. [2, 3]), and coalescence of (non-precessing) binary systems [4].

We distinguish two classes of signal parameters: (i) *four* “amplitude parameters”, namely the amplitudes A_+ and A_\times of the two GW polarizations, the orientation angle ψ of the principal polarization axis, and the initial GW phase ϕ_0 , and (ii) the set of “Doppler parameters” λ , which determine the time evolution of the GW phase $\phi(t; \lambda)$. We restrict our attention to *targeted searches*, in which the Doppler parameters λ are assumed to be known, resulting in a detection problem with four unknown amplitude parameters.

The popular, yet *ad-hoc*, orthodox approach consists of *maximizing* the likelihood function of the data over these four amplitude parameters. It was first shown in [5] that this maximization can be achieved analytically, resulting in a computationally very efficient detection statistic, known as the \mathcal{F} -statistic, which has been used in a number of searches for GWs (e.g. [6, 7, 8, 9]).

Here we investigate an alternative Bayesian approach, which leads us to the Bayes factor as a useful classical detection statistic (see also [10, 11, 12, 13]). Contrary to the maximum-likelihood approach, the Bayesian framework requires the explicit prescription of a prior probability distribution for the unknown signal parameters. We show that a particularly simple, yet unphysical, choice of amplitude prior results in the \mathcal{F} -statistic as a special case of a Bayes factor. This illustrates that frequentist *ad-hoc* statistics often carry their own *unchecked* and *implicit* priors, hidden from view and often unknown to the user (see also [14, 15]).

We can derive a more natural amplitude prior from our model assumption about the emission of GWs from non-axisymmetric spinning systems: the amplitude parameters are closely related to the *orientation* of the emitter with respect to the observer. In the absence of astrophysical information, an isotropic probability distribution for the spin-axis orientation is therefore the natural choice. We refer to the Bayes factor resulting from this amplitude prior as the “ \mathcal{B} -statistic”. Isotropic spin-axis orientation priors have been used previously for Bayesian parameter estimation [6, 16, 17], and in Monte-Carlo simulations to determine frequentist upper limits [7, 18]. A comparison of Bayesian and frequentist methods for setting upper limits can be found in [19, 6].

The \mathcal{F} -statistic (or matched filtering in general) has often been incorrectly referred to as an “optimal statistic” (e.g. [5, 6, 4, 20, 7, 9] to name only a few). Using Monte-Carlo simulations we show that the Bayesian \mathcal{B} -statistic is more powerful (i.e. has a higher expected detection probability at equal false-alarm probability) than the \mathcal{F} -statistic for GWs emitted by systems with random (isotropic) spin-axis orientations. This is a direct consequence of the \mathcal{B} -statistic prior being consistent with the injected distribution of parameters, contrary to the (implicit) \mathcal{F} -statistic prior. Similar results were found previously in the case of burst detection statistics [12, 14, 15].

2. Signal model: coherent gravitational waves of constant amplitudes

The spatial metric perturbation $\vec{h}(t)$ of a coherent GW of constant amplitudes (far from the source) can be written as

$$\vec{h}(t) = \vec{e}_+ A_+ \cos[\phi(t; \lambda) + \phi_0] + \vec{e}_\times A_\times \sin[\phi(t; \lambda) + \phi_0], \quad (1)$$

where $\vec{e}_+ = \hat{u} \otimes \hat{u} - \hat{v} \otimes \hat{v}$, and $\vec{e}_\times = \hat{u} \otimes \hat{v} + \hat{v} \otimes \hat{u}$ are two polarization basis tensors, constructed from a right-handed basis $\{\hat{u}, \hat{v}, -\hat{n}\}$. The unit-vector \hat{n} is pointing along the line of sight from the detector to the source, and the wave-plane basis vectors $\{\hat{u}, \hat{v}\}$ are aligned with the principal polarization axes of the GW. In general the GW phase $\phi(t; \lambda)$ depends on a set of *Doppler parameters* λ , which include the source sky-position \hat{n} , the GW frequency f , and possibly higher-order time derivatives of the frequency $\{\dot{f}, \ddot{f}, \dots\}$. If the source is a neutron star in a binary system, then λ would also include the orbital parameters of the system.

We assume the GW emitter consists of a (non-axisymmetric) rotating system with spin $\vec{\Omega}$ and ellipticity ε with respect to the rotation axis. The corresponding characteristic amplitude h_0 of the GW at the detector can be expressed as

$$h_0 = \frac{4G}{c^4} \frac{I_{zz} \Omega^2}{d} \varepsilon, \quad (2)$$

where I_{zz} is the moment of inertia with respect to the rotation axis, the rotation rate is $\Omega \equiv |\vec{\Omega}|$, and d is the distance to the detector. The emitter *geometry* is




Figure 1. Source geometry angles: ι is the inclination angle between the rotation axis $\vec{\Omega}$ and the line of sight $(-\hat{n})$. The polarization angle ψ measures the orientation of the projected rotation axis $\vec{\Omega}_{\parallel}$ in the sky plane with respect to an observer frame $\{\hat{\xi}, \hat{\zeta}\}$.

fully characterized by two Euler angles describing the orientation of the rotation axis, namely the *inclination angle* ι of the rotation axis with respect to the line of sight \hat{n} , and the *polarization angle* ψ describing the orientation of the projected rotation axis in the sky-plane (see Fig. 1). The two polarization amplitudes A_+ , A_{\times} in (1) can be expressed as

$$A_+ = \frac{1}{2}h_0(1 + \cos^2\iota), \quad \text{and} \quad A_{\times} = h_0 \cos \iota, \quad (3)$$

in terms of the characteristic amplitude h_0 and the inclination angle ι . This corresponds to a choice of the wave-plane coordinate axis \hat{u} perpendicular to the rotation axis [21] and assigns A_+ to the larger of the two principal polarization amplitudes, i.e. $A_+ \geq A_{\times}$. For each sky position \hat{n} we can define a source-independent (right-handed, orthonormal) basis $\{\hat{\xi}, \hat{\zeta}, -\hat{n}\}$, e.g. where $\hat{\xi}$ lies in the ecliptic plane and $\hat{\zeta}$ in the northern hemisphere. This allows us to define the polarization angle ψ as the angle between the principal polarization axis \hat{u} of the GW and the basis vector $\hat{\xi}$, i.e. $\psi \equiv \angle(\hat{u}, \hat{\xi}) = \angle(\hat{v}, \hat{\zeta})$, where $\hat{v} \propto \vec{\Omega}_{\parallel}$ (see Fig. 1). As first shown in [5], the strain $h(t)$ measured in the detector due to a GW signal (1) can be expressed as

$$h(t; \mathcal{A}, \lambda) = \mathcal{A}^{\mu} h_{\mu}(t; \lambda), \quad (4)$$

where we use automatic summation $\sum_{\mu=1}^4$ over repeated amplitude indices μ . The explicit form of the four “basis functions” $h_{\mu}(t; \lambda)$ can be found in [9], for example, but is not important for the following discussion. The four canonical amplitudes \mathcal{A}^{μ} are defined as

$$\begin{aligned} \mathcal{A}^1 &= A_+ \cos 2\psi \cos \phi_0 - A_{\times} \sin 2\psi \sin \phi_0, \\ \mathcal{A}^2 &= A_+ \sin 2\psi \cos \phi_0 + A_{\times} \cos 2\psi \sin \phi_0, \\ \mathcal{A}^3 &= -A_+ \cos 2\psi \sin \phi_0 - A_{\times} \sin 2\psi \cos \phi_0, \\ \mathcal{A}^4 &= -A_+ \sin 2\psi \sin \phi_0 + A_{\times} \cos 2\psi \cos \phi_0. \end{aligned} \quad (5)$$

The set of *amplitude parameters* \mathcal{A} can therefore be expressed either in “physical coordinates” $\bar{\mathcal{A}}^i$, i.e.

$$\{\mathcal{A}\}^i = \bar{\mathcal{A}}^i = \{h_0, \cos \iota, \psi, \phi_0\}, \quad (6)$$

or in “canonical coordinates” \mathcal{A}^{μ} given in (5).

3. Hypothesis testing for GW detection

3.1. Simple versus composite hypotheses

In the following we restrict ourselves to the case where we need to decide only between two hypotheses, namely $\mathcal{H}_N \equiv$ “the data x consist of only noise n ”, and $\mathcal{H}_S \equiv$ “the data contains a signal s in addition to noise n ”, where we assume a signal $s = h(t; \mathcal{A}, \lambda)$ of the form (4). For simplicity we assume the Doppler parameters $\lambda = \lambda_s$ to be known *a-priori*, while the four amplitude parameters $\mathcal{A} \in \mathbb{A}$ are unknown. This corresponds to a *targeted* search for GWs, for example from an isolated pulsar with known sky position and GW frequency $f(t)$. We can formally write the two hypothesis as

$$\begin{aligned} \mathcal{H}_N : x(t) &= n(t), \\ \mathcal{H}_S : x(t) &= n(t) + s(t; \mathcal{A}, \lambda_s), \quad \text{for any } \mathcal{A} \in \mathbb{A}. \end{aligned} \quad (7)$$

Note that \mathcal{H}_N is a *simple* hypothesis, which means that all model parameters are fully specified, namely $s = 0$. On the other hand, \mathcal{H}_S is a *composite* hypothesis, as the amplitude parameters $\mathcal{A} \in \mathbb{A}$ of the signal are unknown. The composite hypothesis \mathcal{H}_S can be considered as a union of simple hypotheses, i.e. $\mathcal{H}_S = \bigcup_{\mathcal{A} \in \mathbb{A}} \mathcal{H}_S(\mathcal{A})$.

3.2. Scalar product and Gaussian noise

Assuming Gaussian stationary noise with known (single-sided) power spectral density S_n , the probability density function (pdf) for a particular pure-noise time series $x = n$ can be written as

$$\text{pdf}(n|S_n) = \kappa e^{-\frac{1}{2}(n|n)}, \quad (8)$$

where κ is a normalization constant and we defined the scalar product $(x|y)$ between time-series $x(t)$ and $y(t)$ as

$$(x|y) \equiv 4\Re \int_0^\infty \frac{\tilde{x}(f) \tilde{y}^*(f)}{S_n(f)} df, \quad (9)$$

where $\tilde{x}(f)$ denotes the Fourier transform of $x(t)$ and $*$ denotes complex conjugation (e.g. see [22]). The likelihood of observing data $x(t)$ in the presence of a signal $s(t)$ is therefore

$$\text{pdf}(x|s S_n) = \kappa e^{-\frac{1}{2}((x-s)|(x-s))}. \quad (10)$$

3.3. Neyman-Pearson optimality

A *detection statistic* $\mathfrak{S}(x)$ is a real-valued function of the data x , such that the corresponding *test* of threshold \mathfrak{S}_* decides for \mathcal{H}_N if $\mathfrak{S}(x) < \mathfrak{S}_*$, and for \mathcal{H}_S if $\mathfrak{S}(x) > \mathfrak{S}_*$. Such a test is typically subject to two types of error: a “false alarm”, i.e. choosing \mathcal{H}_S when \mathcal{H}_N is true, and a “false dismissal”, i.e. rejecting \mathcal{H}_S when it is in fact true. We can express the probability f_A of a false alarm as

$$f_A(\mathfrak{S}_*) = P(\mathfrak{S} > \mathfrak{S}_* | \mathcal{H}_N) = \int_{\mathfrak{S}_*}^\infty \text{pdf}(\mathfrak{S} | \mathcal{H}_N) d\mathfrak{S}. \quad (11)$$

The probability f_D of a false dismissal is $f_D(\mathfrak{S}_*; \Pi_{\mathbb{A}}) = P(\mathfrak{S} < \mathfrak{S}_* | \mathcal{H}_S \Pi_{\mathbb{A}})$, where we notice that the specification (7) of the signal hypothesis \mathcal{H}_S is in fact *incomplete*: in addition to the condition $\mathcal{A} \in \mathbb{A}$, we also need to specify the probability density for \mathcal{A} , which we denote as $\Pi_{\mathbb{A}}$. When measuring f_D in a Monte-Carlo simulation, $\Pi_{\mathbb{A}}$ would

corresponds to the assumed underlying population, from which signals are randomly drawn in each test. The complementary *detection probability* $\eta \equiv 1 - f_D$ is

$$\eta(\mathfrak{S}_*; \Pi_{\mathbb{A}}) = \int_{\mathfrak{S}_*}^{\infty} \text{pdf}(\mathfrak{S} | \mathcal{H}_S \Pi_{\mathbb{A}}) d\mathfrak{S}. \quad (12)$$

Note that this contains the usual definition of the *power function* $\eta(\mathfrak{S}_*; \mathcal{A})$ as a special case, where $\Pi_{\mathbb{A}} = \mathcal{A}$. The definition (12) can also be interpreted as the *expected power* over a population $\Pi_{\mathbb{A}}$. The Neyman-Pearson framework for hypothesis testing defines the *most powerful test* of size f_A as a test that has the highest detection probability η (i.e. smallest false dismissal f_D) for a false-alarm probability of at most f_A (e.g. see [23, 24]). In this framework one compares the respective detection probabilities η of different detection statistics at a given false-alarm probability f_A , which defines the *receiver-operator characteristics* (ROC), namely the function $\eta(f_A; \Pi_{\mathbb{A}})$.

In general the relative performance of different detection statistics will depend on the assumed probability distribution $\Pi_{\mathbb{A}}$ of signal parameters. One statistic can be more efficient in certain parts of parameter space and less efficient in others. If a test is most powerful over the whole parameter space (i.e. has the highest $\eta(\mathfrak{S}_*; \mathcal{A})$ for all $\mathcal{A} \in \mathbb{A}$), it is called a *uniformly most powerful test*.

4. Frequentist maximum-likelihood approach: the \mathcal{F} -statistic

When comparing two *simple* hypotheses, such as \mathcal{H}_N and $\mathcal{H}_S(\mathcal{A})$ for *known* \mathcal{A} , the Neyman-Pearson Lemma states that the most powerful test (cf. section 3.3) is the *likelihood-ratio* \mathcal{L} , defined as

$$\mathcal{L}(x; \mathcal{A}) \equiv \frac{\text{pdf}(x | \mathcal{H}_S(\mathcal{A}))}{\text{pdf}(x | \mathcal{H}_N)}. \quad (13)$$

Assuming Gaussian noise and using (10), we explicitly obtain

$$\mathcal{L}(x; \mathcal{A}) = \exp \left[(x|s) - \frac{1}{2} (s|s) \right]. \quad (14)$$

However, in the case of a *composite* hypothesis \mathcal{H}_S with *unknown* amplitude parameters \mathcal{A} , the orthodox frequentist framework does not generally provide a canonical detection statistic. Interestingly, one cannot even define a frequentist likelihood $\text{pdf}(x | \mathcal{H}_S)$. A common but *ad-hoc* approach to dealing with composite hypotheses consists of using the *maximum* of the likelihood ratio $\mathcal{L}(x; \mathcal{A})$ over the parameter space \mathbb{A} , i.e. define

$$\mathcal{L}_{\text{ML}}(x) \equiv \max_{\mathcal{A} \in \mathbb{A}} \mathcal{L}(x; \mathcal{A}) \quad (15)$$

as a composite-hypothesis test: decide for \mathcal{H}_S if $\mathcal{L}_{\text{ML}}(x) > \mathcal{L}_*$ and \mathcal{H}_N otherwise. Using (14) and (4), the likelihood-ratio function can be written more explicitly as

$$\mathcal{L}(x; \mathcal{A}) = \exp \left[\mathcal{A}^\mu x_\mu - \frac{1}{2} \mathcal{A}^\mu \mathcal{M}_{\mu\nu} \mathcal{A}^\nu \right], \quad (16)$$

where we defined

$$x_\mu \equiv (x | h_\mu(\lambda_s)), \quad \mathcal{M}_{\mu\nu} \equiv (h_\mu(\lambda_s) | h_\nu(\lambda_s)). \quad (17)$$

We see that $\mathcal{L}(x; \mathcal{A})$ is a Gaussian function in \mathcal{A}^μ , so we can analytically maximize it to obtain

$$\mathcal{L}_{\text{ML}}(x) = e^{\mathcal{F}(x)}, \quad \text{with} \quad \mathcal{F}(x) \equiv \frac{1}{2} x_\mu \mathcal{M}^{\mu\nu} x_\nu, \quad (18)$$

with the inverse matrix $\mathcal{M}^{\mu\nu}$ defined via $\mathcal{M}^{\mu\sigma}\mathcal{M}_{\sigma\nu} = \delta_\nu^\mu$. This defines the so-called \mathcal{F} -statistic, which was first derived in this context in [5]. The statistic $2\mathcal{F}(x)$ can be shown to be χ^2 -distributed with four degrees of freedom, and a non-centrality parameter $\rho^2 \equiv (s|s)$, where ρ is called the (optimal) *signal-to-noise ratio*. The expectation value of $2\mathcal{F}$ is $E[2\mathcal{F}] = 4 + \rho^2$.

5. Bayesian hypothesis testing

The Bayesian hypothesis-testing framework follows uniquely from a straightforward application of the probability axioms (cf [25, 26, 10, 27]). For any question of interest one can (at least in principle) *compute* the probability of different hypotheses, optimally using the available information such as the observed data $x(t)$, and all our prior information and assumptions, which we denote by I . Here we use the Bayesian approach to construct a classical *detection statistic*, in order to compare its performance to the frequentist \mathcal{F} -statistic in the Neyman-Pearson framework.

For any hypothesis \mathcal{H}_i we can directly express the probability of \mathcal{H}_i being true given the data x and our background assumptions I , namely

$$P(\mathcal{H}_i|xI) = \frac{\text{pdf}(x|\mathcal{H}_iI)P(\mathcal{H}_i|I)}{\text{pdf}(x|I)}. \quad (19)$$

This expression is known as *Bayes' theorem*, and it follows directly from the product rule of probabilities applied to $\text{pdf}(\mathcal{H}_i|xI)$. The term $P(\mathcal{H}_i|I)$ is the *prior probability* for \mathcal{H}_i . Contrary to the orthodox frequentist framework, the (marginal) likelihood[‡] $\text{pdf}(x|\mathcal{H}_iI)$ of observing data x given \mathcal{H}_i is well-defined even for composite hypotheses. In order to compute $\text{pdf}(x|\mathcal{H}_S I)$, we simply use the product rule to write

$$\text{pdf}(\mathcal{A}|x\mathcal{H}_S I) = \frac{\text{pdf}(\mathcal{A}x|\mathcal{H}_S I)}{\text{pdf}(x|\mathcal{H}_S I)}, \quad (20)$$

and invoking the normalization condition $\int \text{pdf}(\mathcal{A}|x\mathcal{H}_S I) d^4\mathcal{A} = 1$, we obtain

$$\begin{aligned} \text{pdf}(x|\mathcal{H}_S I) &= \int_{\mathbb{A}} \text{pdf}(\mathcal{A}x|\mathcal{H}_S I) d^4\mathcal{A} \\ &= \int_{\mathbb{A}} \text{pdf}(x|\mathcal{A}\mathcal{H}_S I) \text{pdf}(\mathcal{A}|\mathcal{H}_S I) d^4\mathcal{A}. \end{aligned} \quad (21)$$

If $\{\mathcal{H}_i\}_{i=1}^m$ is a set of m mutually exclusive and exhaustive hypotheses, i.e. exactly one of them is true, then one obtains the normalization condition $\sum_{i=1}^m P(\mathcal{H}_i|xI) = 1$, which determines the denominator $\text{pdf}(x|I)$ in (19). We do not need to make this assumption, however, as we can instead compute the *relative* probability of \mathcal{H}_S with respect to \mathcal{H}_N , which is known as the (posterior) *odds ratio* O_{SN} , namely

$$O_{SN}(x|I) \equiv \frac{P(\mathcal{H}_S|xI)}{P(\mathcal{H}_N|xI)} = \frac{\text{pdf}(x|\mathcal{H}_S I)}{\text{pdf}(x|\mathcal{H}_N I)} \frac{P(\mathcal{H}_S|I)}{P(\mathcal{H}_N|I)}. \quad (22)$$

This expressions shows how the prior odds ratio $P(\mathcal{H}_S|I)/P(\mathcal{H}_N|I)$ gets “updated” by the observation of x , namely by multiplication with the (marginal) *likelihood-ratio*

$$B_{SN}(x|I) \equiv \frac{\text{pdf}(x|\mathcal{H}_S I)}{\text{pdf}(x|\mathcal{H}_N I)}, \quad (23)$$

which in also as the *Bayes factor*. Note that the prior odds ratio is a constant factor in O_{SN} , and therefore plays no role in constructing a classical *detection statistic* $\mathfrak{S}(x)$

[‡] also known as the *evidence*

(any monotonic function of $\mathfrak{S}(x)$ has the same power). Using (21) and (13), we can write the Bayes factor (23) explicitly as

$$B_{\text{SN}}(x|I) = \int_{\mathbb{A}} \mathcal{L}(x; \mathcal{A}) \text{pdf}(\mathcal{A}|\mathcal{H}_S I) d^4 \mathcal{A}. \quad (24)$$

Note that while the \mathcal{F} -statistic (18) was obtained by *maximizing* the likelihood ratio $\mathcal{L}(x; \mathcal{A})$ over the “nuisance parameters” \mathcal{A}^μ , the Bayes factor $B_{\text{SN}}(x)$ consists of *marginalizing* $\mathcal{L}(x; \mathcal{A})$ with an amplitude prior $\text{pdf}(\mathcal{A}|\mathcal{H}_S I)$. In order to uniquely specify the Bayes factor, we therefore need to determine the function $\text{pdf}(\mathcal{A}|\mathcal{H}_S I)$, which adequately describes our ignorance of the pulsar amplitude parameters \mathcal{A} .

5.1. Uniform priors in \mathcal{A}^μ -coordinates: rediscovering the \mathcal{F} -statistic

Considering the form (16) of $\mathcal{L}(x; \mathcal{A})$, a straightforward choice would be uniform priors in coordinates \mathcal{A}^μ . We refer to this as the “canonical prior” Π_c , namely

$$\text{pdf}(\{\mathcal{A}^\mu\}|\mathcal{H}_S \Pi_c I) = \begin{cases} C & \text{if } h_0(\{\mathcal{A}^\mu\}) < h_0^{\text{max}}, \\ 0 & \text{otherwise,} \end{cases} \quad (25)$$

where h_0^{max} is the maximum amplitude we consider possible, and $h_0(\{\mathcal{A}^\mu\})$ is given by inversion of Eq. (5). The normalization constant C is determined by $\int_{\mathbb{A}} \text{pdf}(\mathcal{A}|\mathcal{H}_S I) d^4 \mathcal{A} = 1$. The actual choice of h_0^{max} is unimportant for the properties of $B_{\text{SN}}(x)$ as a *detection statistic*, because for large $h_0^{\text{max}} \gg 1$, the marginalization (24) leads to a Gaussian integral, namely

$$\begin{aligned} B_{\text{SN}}(x|\Pi_c I) &= C \int_{\mathbb{A}} \mathcal{L}(x; \mathcal{A}) d^4 \mathcal{A} \\ &= C (2\pi)^2 (\det \mathcal{M})^{-1/2} e^{\mathcal{F}(x)}, \end{aligned} \quad (26)$$

where $\det \mathcal{M}$ is the determinant of the matrix $\mathcal{M}_{\mu\nu}$. We see that uniform amplitude priors in \mathcal{A}^μ -coordinates lead us back to the \mathcal{F} -statistic (18). However, there is an additional antenna-pattern factor $(\det \mathcal{M})^{-1/2}$, which depends on the sky-position \hat{n} and the observation period. For a targeted search with a single known sky position, this is a constant factor which does not affect the power of the detection statistic, i.e. $B_{\text{SN}}(x|\Pi_c I)$ is equivalent to $\mathcal{F}(x)$. This weighting factor would play a role, however, when investigating searches over unknown sky position. A similar effect was first noted in the Bayesian analysis of burst detection statistics [15].

5.2. Physical priors in amplitude-space: introducing the \mathcal{B} -statistic

Despite our assumed “ignorance” about the amplitude parameters of the GW signal, we have made a number of model assumptions about the geometry of the emitting system (see section 2). In a sense the physical model describing the emitting system singles out a preferred coordinate system in \mathbb{A} , in which we should express our ignorance. We refer to the resulting prior as the “physical prior” Π_{ph} .

The initial phase ϕ_0 is directly related to the rotation angle of the quadrupolar deformation with respect to the rotation axis $\hat{\Omega}$ at some fixed reference time. The probability distribution for ϕ_0 is therefore independent of $\{h_0, \cos \iota, \psi\}$, and by rotational symmetry we can assign a uniform prior, i.e.

$$\text{pdf}(\phi_0|\Pi_{\text{ph}} I) = \frac{1}{2\pi}, \quad \phi_0 \in [0, 2\pi). \quad (27)$$

In section 2 and Fig. 1 we have seen that $\cos \iota$ and ψ determine the orientation of the rotation axis $\vec{\Omega}$ with respect to the observer frame $\{\hat{\xi}, \hat{\zeta}, -\hat{n}\}$. If we have no information about the orientation of the emitting system, then rotational symmetry dictates an isotropic probability distribution for $\vec{\Omega}$. The surface element on the unit sphere of $\vec{\Omega}$ orientations can be expressed as $d^2S = |d \cos \iota| |d\psi|$, and because $\cos \iota$ and ψ are independent degrees of freedom, their respective prior probabilities are

$$\text{pdf}(\cos \iota | \Pi_{\text{ph}} I) = \frac{1}{2}, \quad \cos \iota \in [-1, 1], \quad (28)$$

$$\text{pdf}(\psi | \Pi_{\text{ph}} I) = \frac{2}{\pi}, \quad \psi \in [-\pi/4, \pi/4], \quad (29)$$

where we used the fact that gauge transformations $\{\psi \rightarrow \psi + \pi/2, \phi_0 \rightarrow \phi_0 + \pi\}$ leave the observed signal (4) unchanged, so ψ can always be brought into the range $\psi \in [-\pi/4, \pi/4]$. Note that these priors are identical to those used previously in Bayesian parameter estimation [6, 16, 17] and Monte-Carlo simulations for frequentist upper limits [7, 18].

Contrary to the angle variables $\cos \iota, \psi$ and ϕ_0 , there is no unique natural choice of uninformed prior for the amplitude h_0 . One could derive a prior for h_0 from Eq. (2), if astrophysical priors for the deformation ε , spin rate Ω and distance d are available. Other possibilities include a “maximum entropy” prior, or a Jeffreys prior. For simplicity, however, we simply chose a uniform prior, namely

$$\text{pdf}(h_0 | \Pi_{\text{ph}} I) = \frac{1}{h_0^{\text{max}}}, \quad h_0 \in [0, h_0^{\text{max}}]. \quad (30)$$

Combining (27) - (30) we obtain an amplitude prior of the form

$$\text{pdf}(h_0, \cos \iota, \psi, \phi_0 | \Pi_{\text{ph}} I) = \frac{1}{2\pi^2 h_0^{\text{max}}} = C', \quad h_0 \in [0, h_0^{\text{max}}], \quad (31)$$

which for simplicity of notation we refer to as the physical prior, while this qualifier can only be justified for the angle variables. Substituting the prior Π_{ph} in the Bayes factor (24), and assuming $h_0^{\text{max}} \gg 1$, we now obtain

$$B_{\text{SN}}(x | \Pi_{\text{ph}} I) = C' \int_0^\infty dh_0 \int_{-1}^1 d \cos \iota \int_{-\pi/4}^{\pi/4} d\psi \int_0^{2\pi} d\phi_0 \mathcal{L}(x; \mathcal{A}), \quad (32)$$

with the likelihood-ratio $\mathcal{L}(x; \mathcal{A})$ of Eq. (16). We use this Bayes factor as a new classical *detection statistic* $\mathcal{B}(x)$, which we refer to as the “ \mathcal{B} -statistic”, namely

$$\mathcal{B}(x) \equiv B_{\text{SN}}(x | \Pi_{\text{ph}} I). \quad (33)$$

5.3. Relation between amplitude priors Π_c and Π_{ph}

In order to compare the physical amplitude prior Π_{ph} of Eq. (31) to the canonical prior Π_c of Eq. (25), we use the coordinate transformation (5) relating \mathcal{A}^μ and $\bar{\mathcal{A}}^i$. The Jacobian J of this transformation is found as

$$J \equiv \left| \det \left(\frac{\partial \mathcal{A}^\mu}{\partial \bar{\mathcal{A}}^i} \right) \right| = \frac{h_0^3}{4} (1 - \cos^2 \iota)^3. \quad (34)$$

Using the identity

$$\text{pdf}(\{\mathcal{A}^\mu\} | \Pi_c I) d^4 \mathcal{A} = \text{pdf}(\{\bar{\mathcal{A}}^i\} | \Pi_c I) d^4 \bar{\mathcal{A}}, \quad (35)$$

together with the relation $d^4\mathcal{A} = J d^4\bar{\mathcal{A}}$ between volume elements, we can translate Π_c into physical coordinates $\{\bar{\mathcal{A}}^i\} = \{h_0, \cos\iota, \psi, \phi_0\}$, namely

$$\text{pdf}(h_0, \cos\iota, \psi, \phi_0 | \Pi_c I) = \frac{C}{4} h_0^3 (1 - \cos^2\iota)^3, \quad (36)$$

which can be compared to the physical prior Π_{ph} in Eq. (31). We see that Π_c agrees with Π_{ph} in assigning uniform prior probabilities to ϕ_0 and ψ , but the prior densities on $\cos\iota$ and h_0 are very different.

The canonical prior information Π_c , which is *implicit* in the \mathcal{F} -statistic (cf. section 5.1), is therefore found to be rather unphysical: a higher prior probability is assigned to stronger signals compared to weaker ones, and signals with near-linear polarization ($\cos\iota \sim 0$, corresponding to “edge-on” emitters) are given undue weight compared to signals with near-circular polarization ($|\cos\iota| \sim 1$, corresponding to “face-on” emitters). This amounts to postulating a non-isotropic probability distribution for the orientation $\vec{\Omega}$ of spinning GW sources, such that $\vec{\Omega}$ favors orientations orthogonal to the line of sight \hat{n} .

6. Comparing detection efficiencies of $\mathcal{F}(x)$ and $\mathcal{B}(x)$

6.1. Estimating the ROC curves

We use the classical Neyman-Pearson framework (cf. section 3.3) in order to compare the detection efficiency, or “power”, of the \mathcal{F} -statistic (18) and the \mathcal{B} -statistic (32). The practical Monte-Carlo procedure for estimating the ROC curve $\eta(f_A; \Pi_A)$ for any detection statistic \mathfrak{S} is straightforward: first generate a large sample of N_{MC} random draws $\{\mathfrak{S}_N\}$ of the statistic \mathfrak{S} for the case of *no signal*, i.e. $s = 0$. From this distribution we can estimate the false-alarm probability (11) as a function of the threshold \mathfrak{S}_* , namely

$$f_A(\mathfrak{S}_*) \approx \frac{N_{\mathfrak{S}_N > \mathfrak{S}_*}}{N_{\text{MC}}}, \quad (37)$$

where $N_{\mathfrak{S}_N > \mathfrak{S}_*}$ is the number of \mathfrak{S}_N values found above the threshold \mathfrak{S}_* . Similarly, in the *signal case* \mathcal{H}_S , we randomly draw signal parameters from the assumed population Π_A and generate corresponding random draws $\{\mathfrak{S}_S\}$ of the statistic \mathfrak{S} . From this distribution we can estimate the detection probability (12), namely

$$\eta(\mathfrak{S}_*; \Pi_A) \approx \frac{N_{\mathfrak{S}_S > \mathfrak{S}_*}}{N_{\text{MC}}}. \quad (38)$$

Inverting (37) to yield $\mathfrak{S}_*(f_A)$ and substituting this into (38), we obtain the ROC curve $\eta(f_A; \Pi_A)$.

6.2. Parameters used in Monte-Carlo simulation

In targeted searches the Doppler parameters λ of the signal are known, and for simplicity of this example we fixed these parameters as: right ascension $\alpha = 2$ rad, declination $\delta = -0.5$ rad and a constant frequency without spindown. We assume the detector location to be LIGO Hanford, and an observation with GPS start time of $t_0 = 756950413$ and duration of $T = 25$ hours. The resulting numerical components of the antenna-pattern matrix $\mathcal{M}_{\mu\nu}$ of Eq. (17) are found as: $\mathcal{M}_{11} = \mathcal{M}_{33} = \frac{T}{S_n} A$, $\mathcal{M}_{22} = \mathcal{M}_{44} = \frac{T}{S_n} B$, and $\mathcal{M}_{12} = \mathcal{M}_{34} = \frac{T}{S_n} C$, with $A = 0.154$, $B = 0.234$ $C = -0.0104$, with all other components (approximately) zero. These parameters

are given for the sake of completeness, the qualitative conclusions do not depend on these choices. We used $N_{\text{MC}} = 10^6$ random draws for each distribution, and we estimate the errors on $\eta(f_A; \Pi_{\mathbb{A}})$ using a jackknife estimator (see [28]) with 100 subsets. The estimated 1σ errors on the detection probability in the following ROC curves are always less than $\sigma(\eta) < 0.004$.

6.3. Monte-Carlo results

Because \mathcal{H}_S is a composite hypothesis, the ROC curves depend on the choice of injected signal population $\Pi_{\mathbb{A}}$. In order to illustrate the dependency on the amplitude parameter space, we first consider two highly unphysical choices of signal populations, namely (i) $\Pi_{\mathbb{A}}$ consisting of a *single*, linearly polarized signal with $\cos \iota = 0$, $\psi = 0$ and (ii) a *single*, (nearly) circularly polarized signal with $\cos \iota = 0.99$, $\psi = 0$. In both cases we fixed the SNR of the signal to be $\rho = 4$. Note that the choice of ϕ_0 is irrelevant for both \mathcal{F} and \mathcal{B} . These two choices reflect universes in which all spinning GW sources happen to be (i) edge-on or (ii) face-on, without the observer having any knowledge about it. The results of these simulations are shown in Fig. 2. For

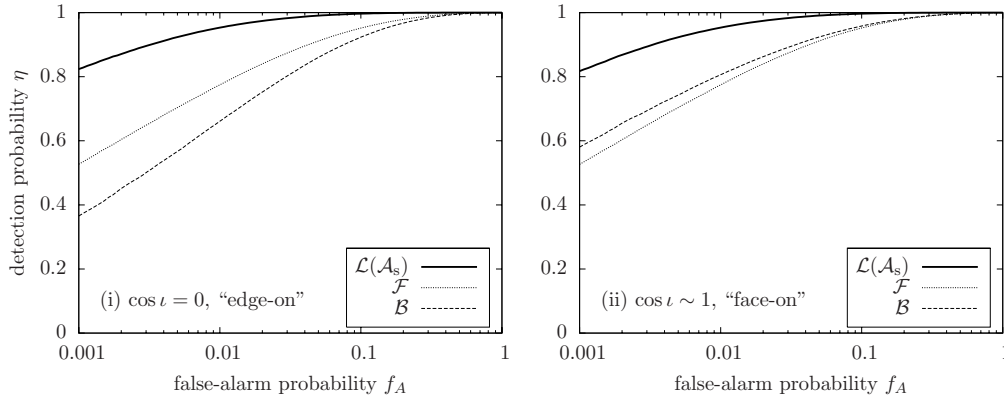


Figure 2. ROC curves $\eta(f_A; \Pi_{\mathbb{A}})$ comparing \mathcal{F} -statistic (18), \mathcal{B} -statistic (33), and the perfect-match likelihood ratio $\mathcal{L}(\mathcal{A}_s)$. The chosen signal populations $\Pi_{\mathbb{A}}$ have fixed SNR of $\rho = 4$, and consist of (i) a *single* linearly polarized signal with $\cos \iota = 0$, $\psi = 0$ (*left panel*), and (ii) a (nearly) circularly polarized signal with $\cos \iota = 0.99$, $\psi = 0$ (*right panel*).

comparison purposes we also plot the perfect-match likelihood ratio $\mathcal{L}(\mathcal{A}_s)$, which is optimal by the Neyman-Pearson lemma for testing the simple hypothesis $\mathcal{H}_S(\mathcal{A}_s)$ (cf. section 3.3), but requires all signal parameters \mathcal{A}_s to be exactly known beforehand.

We see that the \mathcal{F} -statistic is more powerful than the \mathcal{B} -statistic if the signal is linearly-polarized, while the \mathcal{B} -statistic dominates for (near-) circularly-polarized GWs. This is not surprising given that the implicit \mathcal{F} -statistic prior Π_c is biased in favor of linear polarization (cf. section 5.3). The results in Fig. 2 show that neither \mathcal{F} - nor \mathcal{B} -statistic is *uniformly most powerful* (cf. section 3.3) over the amplitude parameter space \mathbb{A} . Note that this does *not* imply that the \mathcal{F} -statistic is more powerful if we *know* a given source to be (near-) linearly polarized. One would fold this knowledge into the prior Π in the Bayes factor (24), while there is no natural way in which this knowledge can be incorporated into the \mathcal{F} -statistic. The resulting Bayes would therefore be more powerful than the \mathcal{F} -statistic.

In the next step we look at more realistic situations in which the injected signals are drawn from a population $\Pi_{\mathbb{A}}$ of randomly distributed $\cos \iota$ and ψ , according to the physical prior Π_{ph} (cf. section 5.2), with (iii) a *fixed SNR* of $\rho = 4$ and (iv) a *fixed amplitude* of $h_0 = 10\sqrt{S_{\text{n}}}$. The results of these simulations are shown in Fig. 3.

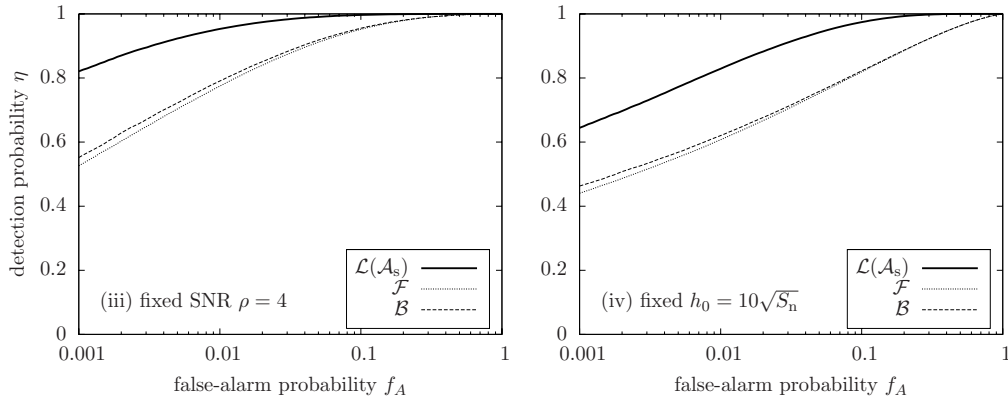


Figure 3. ROC curves $\eta(f_A; \Pi_{\mathbb{A}})$ comparing \mathcal{F} -statistic (18), \mathcal{B} -statistic (33), and the perfect-match likelihood ratio $\mathcal{L}(\mathcal{A}_s)$. The signal populations $\Pi_{\mathbb{A}}$ consist of randomly distributed $\cos \iota$ and ψ , according to the physical prior Π_{ph} (cf. section 5.2), with (iii) a *fixed SNR* of $\rho = 4$ and (iv) a *fixed amplitude* of $h_0 = 10\sqrt{S_{\text{n}}}$.

We see that in these situations the \mathcal{B} -statistic is consistently more powerful than the \mathcal{F} -statistic. This is not surprising, given that the amplitude prior Π_{c} that is implicit in the \mathcal{F} -statistic differs substantially from the injected “real-world” isotropic probability distribution $\Pi_{\mathbb{A}}$ on the orientation of $\vec{\Omega}$. The \mathcal{B} -statistic prior Π_{ph} , on the other hand, is consistent with $\Pi_{\mathbb{A}}$ by construction (cf. section 5.3). In fact, it can be argued [12] that the Bayes factor with a signal prior that is consistent with the population injected in the Monte-Carlo simulation is *by construction* optimal in the sense of the highest expected detection probability at given false-alarm probability. One would therefore not expect any other detection statistic to outperform $\mathcal{B}(x)$ in the simulations shown in Fig. 3.

7. Conclusions

We have shown that the maximum-likelihood \mathcal{F} -statistic can be interpreted as a Bayes factor with a simple, but unphysical, amplitude prior (and an additional unphysical sky-position weighting). Using a more physical prior based on an isotropic probability distribution for the unknown spin-axis orientation of emitting systems, we obtain a new detection statistic, referred to as \mathcal{B} -statistic. Monte-Carlo simulations for signals with random (isotropic) spin-axis orientations show that the \mathcal{B} -statistic is more powerful (in terms of its expected detection probability) than the \mathcal{F} -statistic.

The \mathcal{F} -statistic is therefore not “optimal” in the classical sense. However, the \mathcal{F} -statistic sensitivity appears to be quite comparable to the \mathcal{B} -statistic (see Fig. 3) in the range of parameters considered, while being computationally more efficient (there are no integrations required), and fully characterized by a known simple distribution. Interpreting it as a Bayes factor clarifies its role as a statement about

relative probabilities of hypotheses, and allows one to use the \mathcal{F} -statistic within a fully Bayesian framework. Such a choice would be based on the simplicity and computational efficiency of the \mathcal{F} -statistic, despite the fact that it is not “optimal”.

Acknowledgments

This work has benefited crucially from numerous discussions with colleagues, in particular John T. Whelan, Jolien Creighton, Teviet Creighton, Christian Röver, Holger Pletsch, Graham Woan, Curt Cutler and Chris Messenger.

References

- [1] Prix R (for the LIGO Scientific Collaboration) 2009 *Neutron Stars and Pulsars* ed Becker W (Springer Verlag) p 651
- [2] Królak A, Tinto M and Vallisneri M 2004 *Phys. Rev. D.* **70** 022003 (*Preprint gr-qc/0401108*)
- [3] Whelan J T, Prix R and Khurana D 2008 *Class. Quant. Grav.* **25** 184029 (*Preprint arXiv:0805.1972[gr-qc]*)
- [4] Pai A, Dhurandhar S and Bose S 2001 *Phys. Rev. D.* **64** 042004+ (*Preprint gr-qc/0009078*)
- [5] Jaranowski P, Królak A and Schutz B F 1998 *Phys. Rev. D.* **58** 063001
- [6] Abbott B *et al.* (LIGO Scientific Collaboration) 2004 *Phys. Rev. D.* **69** 082004
- [7] Abbott B *et al.* (LIGO Scientific Collaboration) 2007 *Phys. Rev. D.* **76** 082001 (*Preprint gr-qc/0605028*)
- [8] Astone P *et al.* 2005 *Class. Quant. Grav.* **22** S1243
- [9] Prix R and Whelan J T 2007 *Class. Quant. Grav.* **24** 565
- [10] Veitch J and Vecchio A 2008 *Phys. Rev. D.* **78** 022001+ (*Preprint 0801.4313*)
- [11] Clark J, Heng I S, Pitkin M and Woan G 2007 *Phys. Rev. D.* **76** 043003+ (*Preprint arXiv:gr-qc/0703138*)
- [12] Searle A C 2008 *ArXiv e-prints* (*Preprint 0804.1161*)
- [13] Veitch J and Vecchio A 2008 *Classical and Quantum Gravity* **25** 184010+ (*Preprint 0807.4483*)
- [14] Searle A C, Sutton P J, Tinto M and Woan G 2008 *Class. Quant. Grav.* **25** 114038+ (*Preprint 0712.0196*)
- [15] Searle A C, Sutton P J and Tinto M 2008 *ArXiv e-prints* (*Preprint 0809.2809*)
- [16] Dupuis R J and Woan G 2005 *Phys. Rev. D.* **72** 102002
- [17] Umstätter R, Meyer R, Dupuis R J, Veitch J, Woan G and Christensen N 2004 *Classical and Quantum Gravity* **21** 1655+
- [18] Abbott B *et al.* (LIGO Scientific Collaboration) 2008 *Phys. Rev. D.* **77** 022001
- [19] Gholami I 2007 *Data Analysis of Continuous Gravitational Waves* Ph.D. thesis Albert Einstein Institute
- [20] Krishnan B, Sintes A M, Papa M A, Schutz B F, Frasca S and Palomba C 2004 *Phys. Rev. D.* **70** 082001 (*Preprint arXiv:gr-qc/0407001*)
- [21] Bonazzola S and Gourgoulhon E 1996 *A&A* **312** 675–690
- [22] Finn L S 1992 *Phys. Rev. D.* **46** 5236–5249
- [23] Helstrom C W 1968 *Statistical Theory of Signal Detection* 2nd ed (Oxford: Pergamon Press)
- [24] Stuart A, Ord J K and Arnold S 1999 *Kendall's advanced theory of statistics. Vol.2A: Classical inference and the linear model* (Arnold)
- [25] Jaynes E T 2003 *Probability Theory. The Logic of Science* (Cambridge University Press)
- [26] Sivia D S 1996 *Data Analysis. A Bayesian Tutorial* (Oxford Science Publications)
- [27] Gregory P C 2005 *Bayesian Logical Data Analysis for the Physical Sciences: A Comparative Approach with 'Mathematica' Support* (Cambridge University Press)
- [28] Conway J H and Sloane N J A 1984 *SIAM Journal on Algebraic and Discrete Methods* **5** 294–305
URL <http://link.aip.org/link/?SML/5/294/1>