

Reanalysis of \mathcal{F} -statistic gravitational-wave searches with the higher criticism statistic

M. F. Bennett, A. Melatos

School of Physics, University of Melbourne, Parkville, VIC 3010, Australia

`mfb@unimelb.edu.au`

and

A. Delaigle, P. Hall

Department of Mathematics and Statistics, University of Melbourne, Parkville, VIC 3010, Australia

ABSTRACT

We propose a new method of gravitational wave detection using a modified form of higher criticism, a statistical technique introduced by Donoho & Jin (2004). Higher criticism is designed to detect a group of sparse, weak sources, none of which are strong enough to be reliably estimated or detected individually. We apply higher criticism as a second-pass method to synthetic \mathcal{F} -statistic and \mathcal{C} -statistic data for a monochromatic periodic source in a binary system and quantify the improvement relative to the first-pass methods. We find that higher criticism on \mathcal{C} -statistic data is more sensitive by $\sim 6\%$ than the \mathcal{C} -statistic alone under optimal conditions (i.e. binary orbit known exactly) and the relative advantage increases as the error in the orbital parameters increases. Higher criticism is robust even when the source is not monochromatic (e.g. phase wandering in an accreting system). Applying higher criticism to a phase-wandering source over multiple time intervals gives a $\gtrsim 30\%$ increase in detectability with few assumptions about the frequency evolution. By contrast, in all-sky searches for unknown periodic sources, which are dominated by the brightest source, second-pass higher criticism does not provide any benefits over a first pass search.

Subject headings: gravitational waves — methods: data analysis — methods: statistical — pulsars: general — stars: binaries — stars:neutron

1. Introduction

Direct detection of gravitational waves appears likely in the near future. Existing terrestrial long-baseline interferometers, such as the Laser Interferometer Gravitational-Wave Observatory

(LIGO) and Virgo, have achieved their design sensitivity (Abbott et al. 2009d; Accadia et al. 2012). Next-generation interferometers now under construction are expected to detect tens of events per year, if contemporary estimates of compact binary coalescence rates are correct (Abadie et al. 2010b).

Searches for periodic sources have the advantage of integrating over long observation times to increase the signal-to-noise ratio. The most likely periodic sources detectable by terrestrial interferometers are rapidly rotating, slightly deformed neutron stars (Bildsten 1998; Ushomirsky et al. 2000; Melatos & Payne 2005). They emit at the spin frequency f_* and its first harmonic $2f_*$ (Jaranowski et al. 1998). X-ray timing measurements find $0.3 \text{ kHz} \lesssim f_* \lesssim 0.6 \text{ kHz}$ for low-mass X-ray binaries (Chakrabarty et al. 2003), placing these sources squarely in the LIGO-Virgo band. Directed searches for electromagnetically observed targets (Abbott et al. 2004, 2007b, 2008b; Abadie et al. 2010a, 2011) and blind, all-sky searches for unknown sources (Abbott et al. 2005, 2007a, 2008a, 2009a,b,c) have both been reported.

Targeted searches for known pulsars use the radio ephemeris to guide the search and reduce computational expense by assuming the electromagnetic and gravitational wave phases track each other closely. One search for the Crab pulsar allowed for a mismatch of up to one part in 10^4 between twice the radio pulse frequency and the gravitational wave frequency (Abbott et al. 2008b). Targeted searches for the Crab and Vela pulsars have set direct upper limits on the wave strain of 3.4×10^{-25} and 2.2×10^{-24} respectively, beating the indirect spin-down limits inferred from radio observations (Abbott et al. 2008b; Abadie et al. 2011). A search of 78 pulsars using data from the third and fourth science runs of the LIGO and GEO 600 detectors set upper limits on the wave strain h and ellipticity ε , the tightest of which are $h < 2.6 \times 10^{-25}$ for PSR J1603-7202 and $\varepsilon < 7.1 \times 10^{-7}$ for PSR J2124-3358 (Abbott et al. 2007b).

Blind, all-sky searches address more sources than targeted searches but are expensive computationally as they cover a larger parameter domain to keep track of the unknown frequency evolution. A number of LIGO all-sky searches for periodic sources have been conducted (Abbott et al. 2005, 2007a, 2008a, 2009a,b,c), some of which leverage the distributed computing power of the Einstein@Home project (Abbott et al. 2009a,c). Frequently, they are based on a maximum-likelihood detection statistic, called the \mathcal{F} -statistic (Jaranowski et al. 1998). The \mathcal{F} -statistic is computed by coherently integrating over the observation time T_{obs} , assuming a biaxial neutron star at a specific spin frequency and sky position. Cutler & Schutz (2005) generalized the \mathcal{F} -statistic to apply to multiple detectors and sources.

The computational expense of a fully coherent search for unknown sources becomes prohibitive as T_{obs} increases. Semi-coherent methods address this problem, dividing the observation time into intervals, which are individually searched coherently but combined incoherently. Semi-coherent methods are more sensitive than a fully coherent search for the same computational cost if the parameter space is large (Wette 2012). Abbott et al. (2008a) reported results of a semi-coherent all-sky search for periodic sources. They described and compared three semi-coherent methods:

StackSlide, PowerFlux, and Hough number count (Brady & Creighton 2000; Krishnan et al. 2004; Cutler et al. 2005; Abbott et al. 2008a). For binary sources, Messenger & Woan (2007) derived the \mathcal{C} -statistic, which incoherently combines individual \mathcal{F} -statistic values with a comb template that matches the orbital sideband structure.

In this paper, we propose a new method for enhancing gravitational wave searches using a second-pass method known as higher criticism. Higher criticism was suggested originally by John Tukey and developed further by Donoho & Jin (2004), who proved that it has some mathematical optimality properties in the case of a specific noise distribution. The ultimate goal of any gravitational wave search is to identify an individual object as a definite source. Higher criticism is not suited for this purpose. It is designed to search against a known background for a group of sparse signals too weak to be detected individually. Higher criticism detects the presence of the group but cannot reliably estimate the waveform. We discuss applications of higher criticism to gravitational wave detection, focusing on the \mathcal{F} -statistic (Jaranski et al. 1998) as a case study. In general, we find that higher criticism is significantly more robust than other methods; it ignores some information that other detection statistics incorporate, so may be less sensitive, but makes fewer assumptions about the form of the signal, especially important where the signal evolves during the search.

The paper is structured as follows. In Section 2, we define the higher criticism statistic, compare the detectability of signals with different amplitudes and sparsities, and compute detection thresholds. In Section 3, we briefly review the \mathcal{F} -statistic. In Section 4, we apply a form of higher criticism to targeted searches for a binary pulsar and compare its performance with the \mathcal{C} -statistic. We discuss how higher criticism can accommodate phase wandering in Section 5. While we apply higher criticism specifically to periodic gravitational wave searches in this paper, it is a general method with other applications, including detecting non-Gaussianity in Wilkinson Microwave Anisotropy Probe (WMAP) data (Cayón et al. 2005).

2. Higher Criticism

2.1. Theory

The higher criticism statistic was introduced by Donoho & Jin (2004) to test the hypothesis that n independent and identically distributed (i.i.d.) samples X_1, \dots, X_n come from the same zero-mean Gaussian distribution $N(0, 1)$, against the alternative that a small fraction of them have a nonzero mean $\mu > 0$, that is, to test:

$$H_0 : X_i \stackrel{\text{i.i.d.}}{\sim} N(0, 1) \quad \text{for } i = 1, \dots, n, \tag{1}$$

$$H_1 : X_i \stackrel{\text{i.i.d.}}{\sim} (1 - \epsilon)N(0, 1) + \epsilon N(\mu, 1) \quad \text{for } i = 1, \dots, n. \tag{2}$$

For example, H_0 could correspond to a situation where all samples are just background noise with no signal ($\mu = 0$), whereas in H_1 a fraction ϵ of samples contain signal ($\mu \neq 0$). Higher criticism is

particularly designed for the very challenging situation where the signal is sparse ($\epsilon \ll 1$) and weak (μ is so small that the largest extreme values under H_1 are essentially the same as those under H_0). Although we introduce the statistic in the context of Gaussian noise, it can be generalized easily to other distributions (see Section 2.3). The case of correlated, or non-white, noise has been treated by Hall & Jin (2010).

Let Z denote a random variable with distribution $N(0, 1)$. The higher criticism statistic HC is computed from the n p-values $p_i = P[Z > X_i]$ generated by the n tests $H_{0,i} : X_i \stackrel{\text{i.i.d.}}{\sim} N(0, 1)$ against $H_{1,i} : X_i \stackrel{\text{i.i.d.}}{\sim} N(\mu, 1)$, $\mu > 0$, $i = 1, \dots, n$. It is defined by (Donoho & Jin 2004),

$$HC = \max_{1 \leq i \leq n} \frac{\sqrt{n}[i/n - p_{(i)}]}{\sqrt{p_{(i)}[1 - p_{(i)}]}}, \quad (3)$$

where $p_{(1)} \leq p_{(2)} \leq \dots \leq p_{(n)}$ are the p_i 's sorted into increasing order. The statistic HC rejects H_0 when $HC > g(n, \alpha)$, where $g(n, \alpha)$ is a threshold chosen so that $P_{H_0}[HC > g(n, \alpha)] = P_{H_0}(\text{reject } H_0) \leq \alpha$, where P_{H_0} denotes the probability when H_0 is true. [Our notation differs from Donoho & Jin (2004) in that we use the symbol g for the threshold instead of h to avoid confusion with the gravitational wave strain h .] In other words, the higher criticism test detects a signal when $HC > g(n, \alpha)$. Donoho & Jin (2004) showed that, when H_0 is true, one has $HC \sim \sqrt{2 \log \log(n)}$ asymptotically as $n \rightarrow \infty$ for all $\alpha > 0$, and that if $0 < \alpha < 1$ and n is large enough, one can use the threshold $g(n, \alpha) \approx \sqrt{2 \log \log(n)}$. However for more general α and n , this does not work well in practice, as we illustrate in Section 2.4, where we discuss how to choose $g(n, \alpha)$ accurately as a function of n and α for finite n .

Next, we quantify more precisely how small ϵ and μ can be for the test based on HC to be able to distinguish H_0 from H_1 . Let r and β be two positive parameters. Donoho & Jin (2004) study the properties of the HC statistic for μ and ϵ of the form

$$\mu = \sqrt{2r \log n}, \quad (4)$$

$$\epsilon = n^{-\beta}. \quad (5)$$

These represent very difficult situations since, with r small, μ of this form is smaller than the mean of the upper extreme statistics, and with β large, the proportion of samples with non zero mean is very close to zero. Clearly, a detection is not possible if r is very small and β is very large at the same time. Donoho & Jin (2004) showed that, as $n \rightarrow \infty$, there exists a function $\rho(\beta)$ defined by

$$\rho(\beta) = \begin{cases} \beta - 1/2, & 1/2 < \beta \leq 3/4, \\ (1 - \sqrt{1 - \beta})^2, & 3/4 < \beta < 1, \end{cases} \quad (6)$$

such that H_0 and H_1 are distinguishable only if $r > \rho(\beta)$. If $r < \rho(\beta)$, there does not exist a test that can distinguish H_0 from H_1 . Donoho & Jin (2004) proved that, under some conditions, the HC statistic is able to distinguish H_0 from H_1 throughout the detectable region $r > \rho(\beta)$, and that it has full asymptotic power, that is, $P_{H_1}(\text{reject } H_0) \rightarrow 1$ as $n \rightarrow \infty$ (here P_{H_1} denotes the

conditional probability when H_1 is true). Another test has also been considered in the literature, namely the Neyman-Pearson likelihood ratio test, but it is less attractive since it requires r and β to be known (Donoho & Jin 2004), whereas the test based on the higher criticism statistic does not. Interestingly, Donoho & Jin (2004) also show that, in addition to the detection boundary, there is a second region $r > \beta$ called the estimable region, where H_0 and H_1 can be distinguished and the mean μ can also be estimated consistently. Figure 1, which can be found in Donoho & Jin (2004), illustrates the detectable and estimable regions on the r - β plane.

2.2. Distribution of HC

In this section, we present briefly some results from Monte-Carlo simulations to illustrate the behavior of HC as a function of signal amplitude μ and sparsity ϵ for $n = 10^6$ sample values drawn randomly from the H_1 distribution given by equation (2).

Figure 2 displays a histogram of HC values for 10^3 trials with $\epsilon = 5 \times 10^{-3}$ and $\mu = 0, 0.1, 0.3, 1$. Figure 3 displays a similar histogram with $\mu = 1$ and $\epsilon = 0, 5 \times 10^{-4}, 2 \times 10^{-3}, 5 \times 10^{-3}$. The top left panel in both figures consists of samples drawn from the null distribution H_0 . As μ or ϵ increases, so too does the average value of HC .

2.3. HC for a general noise distribution

In Section 2.1, we defined HC in the context of Gaussian noise. As shown in Donoho & Jin (2004), higher criticism can be generalized to other noise and signal distributions, \mathcal{N} and $\mathcal{S}(\mu)$ say, where μ describes the amplitude of the signal (e.g. mean of the Gaussian distribution in Section 2.1). In this more general setting we test to discriminate between the hypotheses

$$H_0 : X_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N} \quad \text{for } i = 1, \dots, n, \quad (7)$$

$$H_1 : X_i \stackrel{\text{i.i.d.}}{\sim} (1 - \epsilon)\mathcal{N} + \epsilon\mathcal{S}(\mu) \quad \text{for } i = 1, \dots, n. \quad (8)$$

In this context, the higher criticism statistic is computed as in equation (3), but replacing the $p_{(i)}$'s by the ordered values of $p_i = P[Z > X_i]$, $i = 1, \dots, n$, where Z is a random variable that has the distribution \mathcal{N} ; see also Delaigle & Hall (2009).

The detection boundary varies with the noise distribution. For example, when \mathcal{N} is the χ^2 distribution that we use later, the proportion ϵ is defined as in equation (5), but the intensity of the signal is defined in terms of the noncentrality parameter ρ^2 of the χ^2 distribution, through

$$\rho^2 = 2r \log n \quad (9)$$

instead of the mean μ in equation (4). With these definitions of r and β , the detection boundary is given by equation (6), as in the normal case.

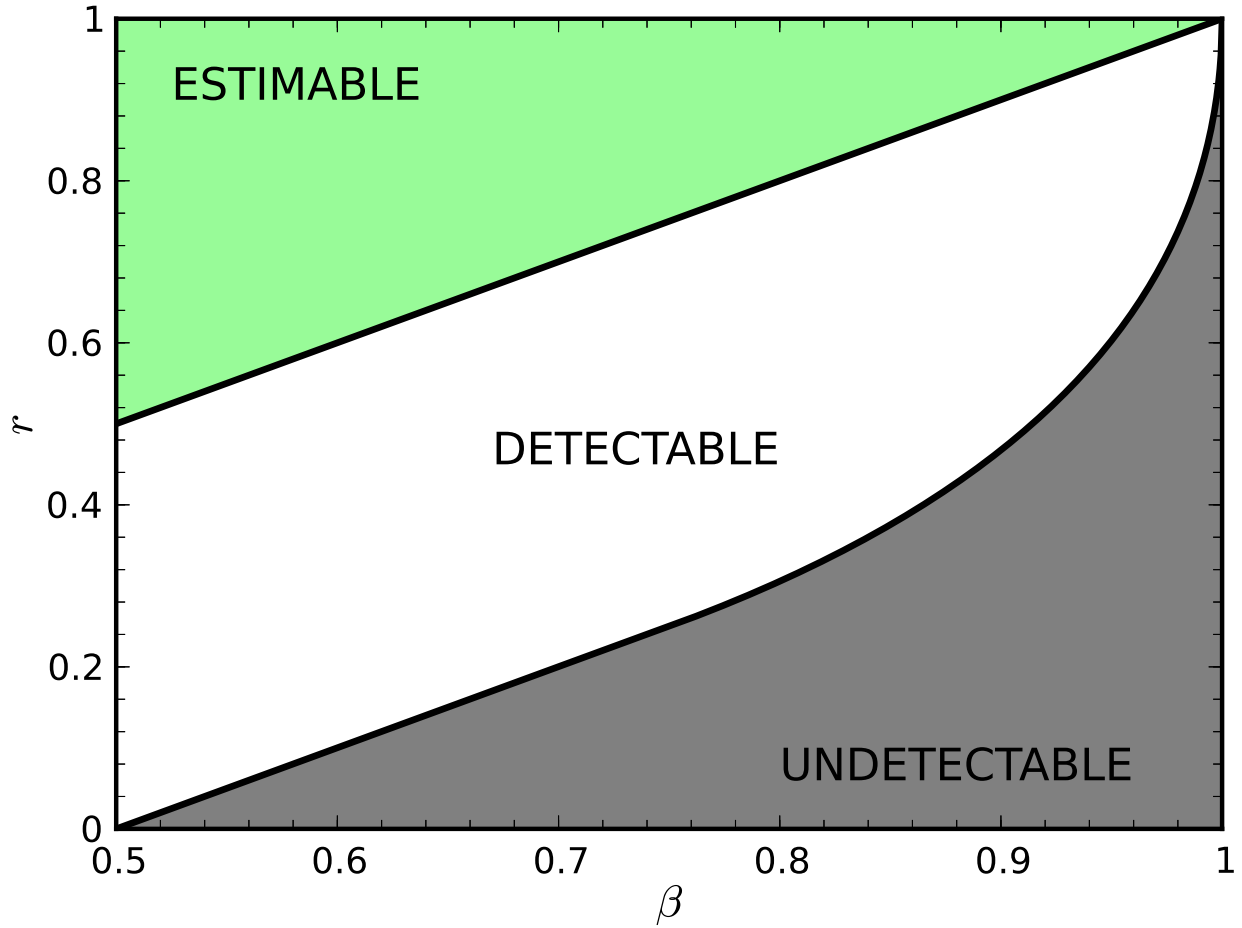


Fig. 1.— Detectability regions on the r - β plane. Undetectable region, shaded gray; detectable region, where signal can be detected but not estimated, unshaded; estimable region, where signal can be both detected and estimated, shaded green.




Fig. 2.— Histograms of HC for 10^3 Monte-Carlo trials, each constructed from $n = 10^6$ samples drawn from distribution H_1 with sparsity $\epsilon = 5 \times 10^{-3}$ and increasing amplitude μ . Upper left panel: noise only ($\mu = 0$), i.e. H_0 distribution. Remaining panels: $\mu = 0.1, 0.3, 1$ (upper right, lower left, lower right).




Fig. 3.— Histograms of HC for 10^3 Monte-Carlo trials, each constructed from $n = 10^6$ samples drawn from distribution H_1 with amplitude $\mu = 1$ and increasing sparsity ϵ . Upper left panel: noise only ($\epsilon = 0$), i.e. H_0 distribution. Remaining panels: $\epsilon = 5 \times 10^{-4}, 2 \times 10^{-3}, 5 \times 10^{-3}$ (upper right, lower left, lower right).

2.4. Detection thresholds

In this section, we discuss the choice of the detection threshold $g(n, \alpha)$ for HC . As already noted earlier, the threshold satisfies $P_{H_0}[HC > g(n, \alpha)] = P_{H_0}(\text{reject } H_0) \leq \alpha$, and under somewhat restrictive conditions, it can be approximated asymptotically by $\sqrt{2 \log \log(n)}$. However, in finite samples and under more general conditions, this theoretical value $\sqrt{2 \log \log(n)}$ is not a good approximation to $g(n, \alpha)$.

To illustrate this we conduct Monte Carlo simulations to determine the threshold $g(n, \alpha)$ for finite n . Note that, when H_0 is true, the p-values are always independent and identically distributed according to the uniform distribution $U(0, 1)$, and thus $g(n, \alpha)$ is independent of the specific noise distribution \mathcal{N} . Therefore it suffices to generate p-values from the $U(0, 1)$ distribution to compute the threshold. We run 10^6 noise-only simulations with $n = 10^3, 10^4, 10^5, 10^6$ to determine Monte-Carlo thresholds. For each, we generate n p-values from $U(0, 1)$ and compute HC .

Figure 4 displays the cumulative distribution function of HC for the simulated noise samples. Threshold values obtained from the simulations by solving an empirical version of $P_{H_0}(HC > g(n, \alpha)) \approx \alpha$ are listed in Table 1 for $n = 10^3, 10^4, 10^5, 10^6$ and $\alpha = 0.5, 0.1, 0.05, 0.01$. For comparison, the theoretical value $\sqrt{2 \log \log(n)}$ is also quoted. From this and Figure 4, we can see that this theoretical value generally does not work well in practice. In each case, a large fraction ($\gtrsim 50\%$) of noise-only trials fall above this threshold. In the remainder of this paper, we commonly use the threshold $g(n, 0.01)$ as determined from Monte-Carlo simulations with $\alpha = 0.01$ (1% false alarm rate).

3. \mathcal{F} -statistic

The \mathcal{F} -statistic is an efficient detection statistic for periodic gravitational waves based on maximum likelihood. A number of blind and targeted searches for periodic sources have been carried out using it and its relatives (Abbott et al. 2004, 2005, 2007a, 2008a,b, 2009a,c; Abadie

Table 1: HC threshold $g(n, \alpha)$.

False alarm rate α	HC threshold $g(n, \alpha)$			
	$n = 10^3$	$n = 10^4$	$n = 10^5$	$n = 10^6$
0.5	2.10	2.26	2.37	2.46
0.1	3.62	3.66	3.70	3.73
0.05	4.72	4.73	4.73	4.73
0.01	10.0	10.1	10.1	10.0
Theoretical threshold $\sqrt{2 \log \log(n)}$	1.97	2.11	2.21	2.29




Fig. 4.— Cumulative distribution function of HC when $n = 10^3$ (red), $n = 10^4$ (blue), $n = 10^5$ (green), and $n = 10^6$ (black) samples drawn from H_0 . Dashed vertical lines indicate the corresponding theoretical $n \rightarrow \infty$ threshold $g(n, \alpha) = \sqrt{2 \log \log(n)}$.

et al. 2010a, 2011). We use the notation $\chi_k^2(\lambda)$ to represent the noncentral (central if $\lambda = 0$) χ^2 distribution with k degrees of freedom and noncentrality parameter λ . In the absence of a signal, and assuming stationary, Gaussian noise, $2\mathcal{F}$ is distributed according to a $\chi_4^2(0)$ distribution, while in the presence of a signal, $2\mathcal{F}$ obeys a $\chi_4^2(\rho^2)$ distribution (Jaranowski et al. 1998; Abbott et al. 2007a). In terms of the waveform, one writes

$$\rho^2 = \frac{2}{S_h(f)} \int_0^{T_{obs}} dt h(t)^2, \quad (10)$$

where $S_h(f)$ is the spectral noise density of the interferometer, T_{obs} is the total observation time, and $h(t) = F_+(t)h_+(t) + F_\times(t)h_\times(t)$ is the gravitational wave strain at the detector, which depends on the beam pattern functions F_+ and F_\times . For a biaxial star, the strains $h_+(t)$ and $h_\times(t)$ depend on the wobble angle θ between the rotation and symmetry axes, the inclination angle ι between the rotation axis and line of sight, right ascension α_{sky} , declination δ_{sky} , polarization angle ψ , and the characteristic strain amplitude

$$h_0 = \frac{16\pi^2 G}{c^4} \frac{\varepsilon I f^2}{D}, \quad (11)$$

where ε is the ellipticity, I is the moment of inertia, and D is the distance to the star (Jaranowski et al. 1998). We obtain an approximate relation between ρ^2 and h_0 by averaging equation (10) over all relevant angles (Jaranowski et al. 1998; Vigelius & Melatos 2009),

$$\rho^2 \approx \frac{32}{375} \frac{h_0^2 T_{obs}}{S_h}, \quad (12)$$

with the assumption that the spectral noise density is achromatic, viz. $S_h(f) = S_h$. Equation (12) allows us to translate detection limits on ρ^2 into limits on h_0 throughout the rest of the paper.

For a known source, the amplitude of the average signal that can be detected coherently with a 1% false alarm rate and a 10% false dismissal rate is (Wette 2012)

$$h_0 = 15.6 \sqrt{\frac{S_h(f)}{T_{obs}}}. \quad (13)$$

Equation (13) is often seen with a factor of 11.4 instead of 15.6 (e.g., Abbott et al. 2007a). The 11.4 factor is obtained assuming the wobble angle $\theta = \pi/2$, whereas we have averaged over θ to arrive at equation (13). For an unknown source, we search across a large number of bins and h_0 must be larger than the value in equation (13) to rise above the background.

Cutler & Schutz (2005) derived the \mathcal{F} -statistic for multiple detectors or multiple sources and found that combining the \mathcal{F} -statistic values from multiple sources increases their detectability as long as the sources have a squared signal-to-noise ratio greater than one fifth of the brightest source. For example, of the known millisecond pulsars, it is advantageous to combine the five strongest ones but no more.

Semi-coherent searches divide the observation time into N intervals of length $\Delta T = T_{obs}/N$. The \mathcal{F} -statistic is computed for each time interval i and combined to obtain

$$2\mathcal{F}_{sc} = \sum_{i=1}^N 2\mathcal{F}_i . \quad (14)$$

As each $2\mathcal{F}_i$ follows a χ_4^2 distribution, $2\mathcal{F}_{sc}$ follows a χ_{4N}^2 distribution.

Searches for periodic sources require detailed knowledge of the frequency evolution. For pulsars, the radio ephemeris is used to guide the search. Any difference between the radio and gravitational wave phases causes significant problems. For unknown sources, we have even less information and are powerless to correct for phase wandering. Recently, Cutler (2011) proposed a “phase-relaxed” \mathcal{F} -statistic for an all-sky search, modifying the semi-coherent detection statistic to accommodate a phase offset δ_i for each coherent time interval i , while preserving a monochromatic phase model overall.

In this paper we outline a new approach: applying higher criticism on a second pass to reanalyze the results of gravitational wave searches. Higher criticism relies on detailed knowledge of the background noise distribution. The detection statistics discussed so far assume stationary, Gaussian noise. Realistic detector noise is neither stationary nor Gaussian. We discuss this key point further in Section 6. For now, we persevere with the common, simplifying assumption of stationary, Gaussian detector noise in order to assess the overall viability of higher criticism, noting that modified forms of higher criticism have been constructed for correlated noise (Hall & Jin 2010) or when the noise distribution is imperfectly known (Delaigle et al. 2011).

4. Targeted binary search

Low-mass X-ray binaries (LMXBs) are accreting neutron stars in a binary orbit with a low mass companion. X-ray pulsations and burst oscillations place the spin frequencies of discovered objects in the range $270 \text{ Hz} \leq f_* \leq 619 \text{ Hz}$ (Chakrabarty et al. 2003), well below the centrifugal break-up frequency (Cook et al. 1994). Gravitational radiation is one effective means to balance the accretion torque and stall spin up (Bildsten 1998). In this section, we compare a second-pass search with higher criticism to the side-band \mathcal{C} -statistic algorithm proposed by Messenger & Woan (2007).

4.1. \mathcal{F} -statistic and \mathcal{C} -statistic

The gravitational wave signal from a neutron star in a binary has its frequency modulated into sidebands with spacing $1/P$, where P is the orbital period (Messenger & Woan 2007). Therefore, an \mathcal{F} -statistic search sees a large number of relatively weak signals across many frequency bands.

For a neutron star with intrinsic gravitational wave frequency f_0 , $2\mathcal{F}$ obeys a $\chi_4^2[\rho^2(f)]$ distribution, where the frequency dependent noncentrality parameter can be written (Messenger & Woan 2007)

$$\rho^2(f) = \rho_0^2 \sum_{n=-\lfloor Z_0 \rfloor}^{\lfloor Z_0 \rfloor} J_n^2(Z_0) \left| \tilde{W}(f - f_{0,n}) \right|^2, \quad (15)$$

with $Z_0 = 2\pi f_0 a$, where $\lfloor Z_0 \rfloor$ represents the largest integer less than Z_0 , J_n is the n -th order Bessel function of the first kind, a is the light crossing time of the orbital radius projected onto the line of sight, $f_{0,n} = f_0 - n/P$ is the frequency of the n th sideband around a source with frequency f_0 , ρ_0^2 is the source signal-to-noise ratio squared [see equation (10)], and $\tilde{W}(f)$ is the Fourier transform of the window function $W(t)$, which equals 1 or 0 when the detector is on or off respectively. Assuming no gaps in the data, one finds

$$\left| \tilde{W}(f) \right|^2 = \frac{\sin^2(\pi f T_{obs})}{\pi^2 f^2 T_{obs}^2}. \quad (16)$$

With the window function defined as in Messenger & Woan (2007), an extra factor of $1/T_{obs}^2$ is included in equation (16) so that \mathcal{C} is dimensionless.

The \mathcal{C} -statistic sums incoherently the \mathcal{F} -statistic power in the orbital sidebands (Messenger & Woan 2007) as follows:

$$\mathcal{C}(f) = \sum_{n=-\lfloor Z_0 \rfloor}^{\lfloor Z_0 \rfloor} 2\mathcal{F}(f - n/P). \quad (17)$$

For noise only, $2\mathcal{F}$ is distributed according to a χ_4^2 distribution. Hence, \mathcal{C} , which is the sum of M $2\mathcal{F}$ values, obeys a χ_{4M}^2 distribution, where $M = 2\lfloor Z_0 \rfloor + 1$ is the total number of sidebands (Messenger & Woan 2007). In the presence of a signal, $2\mathcal{F}(f)$ is distributed according to a noncentral $\chi_4^2[\rho^2(f)]$ distribution with a frequency dependent noncentrality parameter $\rho^2(f)$ and \mathcal{C} follows a noncentral $\chi_{4M}^2(\lambda)$ distribution with noncentrality parameter (Messenger & Woan 2007)

$$\lambda = \sum_{n=-\lfloor Z_0 \rfloor}^{\lfloor Z_0 \rfloor} \rho^2(f - n/P). \quad (18)$$

4.2. Higher criticism

How does a \mathcal{C} -statistic search, which assumes that P (and hence the sideband locations) are known, perform in comparison with a second-pass higher criticism search over \mathcal{F} -statistic or \mathcal{C} -statistic values? We run Monte-Carlo simulations to answer this question.

We simulate a targeted search for Sco X-1, similar to the one described in Sammut et al. (2013, in preparation). The search parameters are $P = 68023.84$ s, $a = 1.44$ s, $T_{obs} = 10$ days, and source frequency $f_0 = 400$ Hz (assumed). We search over the range of frequencies $100 \leq f \leq 1000$ Hz with frequency bin spacing $\delta f = 1/(2T_{obs}) = 5.8 \times 10^{-7}$ Hz, corresponding to $n = 1.5 \times 10^9$ frequency

bins in total. The signal is modulated into M sidebands separated by $1/P$. M depends on the signal frequency, with $1811 \leq M \leq 18097$ for $100 \text{ Hz} \leq f_0 \leq 1000 \text{ Hz}$. The width in frequency space of the entire sideband structure (the comb) is $0.03 \text{ Hz} \leq M/P \leq 0.27 \text{ Hz}$.

Let us apply higher criticism to the problem. We take advantage of the relatively narrow sideband structure and divide the frequency domain $[f_{min}, f_{max}]$ into N_w windows of equal width w , each containing n_w frequency bins, viz. $[f_{min}, f_{min} + w]$, $[f_{min} + w/2, f_{min} + 3w/2]$, $[f_{min} + w, f_{min} + 2w]$, ..., $[f_{max} - 3w/2, f_{max} - w/2]$, $[f_{max} - w, f_{max}]$. The width w is chosen to be twice the maximum width of the sidebands, $w = 2M/P$, and windows overlap by 50% to ensure that the entire signal is contained entirely within a single window. A more sophisticated method would vary the width with frequency as the sideband width increases from 0.03 Hz for $f_0 = 100 \text{ Hz}$ to 0.27 Hz for $f_0 = 1000 \text{ Hz}$, however, for simplicity, we construct our search windows based on the maximum sideband width at $f_{max} = 1000 \text{ Hz}$, obtaining $w = 0.53 \text{ Hz}$, $n_w = 9.2 \times 10^5$, and $N_w = 3.4 \times 10^3$.

The data are synthesized as follows. First, the noncentrality parameter is computed for each frequency bin. The signal is modulated into sidebands and the frequency dependent noncentrality can be written (Sammut et al. 2013, in preparation) ,

$$\rho^2(f) \approx \rho_0^2 \sum_{n=-\lfloor Z_0 \rfloor}^{\lfloor Z_0 \rfloor} J_n^2(Z_0) \left| \tilde{W}(f - f_{0,n}) \right|^2 . \quad (19)$$

We simplify this equation assuming $\left| \tilde{W}(f) \right|^2 \approx \delta(f)$ (true in the limit $T_{obs} \rightarrow \infty$) and assume the $J_n^2(Z_0)$ all have similar amplitude, which we absorb into the constant ρ_0^2 . Additionally, we adjust the frequency of each sideband $f_{0,n}$ to be equal to the corresponding closest frequency bin, $\text{nr}(f_{0,n})$, using $\text{nr}(f)$ to represent the frequency bin closest to frequency f , neglecting the potential loss of amplitude caused by the mismatch between $f_{0,n}$ and $\text{nr}(f_{0,n})$. We denote the frequency on the i -th frequency bin f'_i , using the ‘prime’ notation to differentiate between frequency bins and the exact frequency of the signal f_0 or the sidebands $f_{0,n}$. The simplified noncentrality parameter used to generate the \mathcal{F} -statistic for each frequency bin f'_i is

$$\rho^2(f'_i) = \rho_0^2 \sum_{n=-\lfloor Z_0 \rfloor}^{\lfloor Z_0 \rfloor} K [f'_i, \text{nr}(f_{0,n})] , \quad (20)$$

where $K(x, y) = \delta_{x,y}$ is the standard Kronecker delta rebranded so that we can see the subscripts clearly. Equation (20) corresponds to $\rho^2(f'_i)$ equal to ρ_0^2 for the frequency bin closest to each of the true sideband locations and zero elsewhere. We generate synthetic $2\mathcal{F}$ values drawn from the $\chi_4^2[\rho^2(f'_i)]$ distribution, for each frequency bin f'_i .

To compute \mathcal{C} -statistic values, we construct a comb centered at each frequency bin f'_i assuming $f_0 = f'_i$,

$$q(f'_i) = \sum_{n=-\lfloor Z'_i \rfloor}^{\lfloor Z'_i \rfloor} K [f'_i, \text{nr}(f'_{i,n})] , \quad (21)$$

with $Z'_i = 2\pi f'_i a$ and $f'_{i,n} = f'_i + n/P$. The \mathcal{C} -statistic is computed by convolving the $2\mathcal{F}(f)$ with $q(f)$ (Messenger & Woan 2007)

$$\mathcal{C}(f'_i) = 2\mathcal{F} * q(f'_i) . \quad (22)$$

We compute the higher criticism statistic as described in Section 2.3 for the $2\mathcal{F}$ and the \mathcal{C} values contained in a given search window and call the results $HC_{2\mathcal{F}}$ and $HC_{\mathcal{C}}$ respectively. The noise distribution for $2\mathcal{F}$ is $\chi_4^2(0)$. It is more complicated for \mathcal{C} , as the number of sidebands $M'_i = 2\lfloor Z'_i \rfloor + 1$ depends on the frequency f'_i . Hence, the p-values for each $\mathcal{C}(f'_i)$ are calculated with respect to the $\chi_{4M'_i}^2$ noise distribution (four degrees of freedom for each sideband summed).

To compare the performance of the three detection statistics \mathcal{C} , $HC_{2\mathcal{F}}$ and $HC_{\mathcal{C}}$, we compute thresholds with a false alarm rate $\alpha = 0.01$. We have n $2\mathcal{F}$ and \mathcal{C} values, one for each frequency bin, and N_w $HC_{2\mathcal{F}}$ and $HC_{\mathcal{C}}$ values, one for each search window. Thresholds are found by insisting that noise-only data fall below the threshold in every bin or window with probability at least $1 - \alpha$ [e.g., Wette (2012)],

$$\{P[\mathcal{C} < \mathcal{C}_{th}(\alpha)]\}^n \geq 1 - \alpha , \quad (23)$$

$$\{P[HC_{2\mathcal{F},\mathcal{C}} < g(n_w, \alpha)]\}^{N_w} \geq 1 - \alpha , \quad (24)$$

where $\mathcal{C}_{th}(\alpha)$ is the \mathcal{C} -statistic threshold. The HC threshold $g(n_w, \alpha)$ is determined from Monte-Carlo simulations, as described in Section 2.4. We integrate the \mathcal{C} -statistic probability distribution to find $\mathcal{C}_{th}(\alpha)$ and use the result to compute the threshold noncentrality parameter $\lambda_{th}(\alpha, \delta)$ for false dismissal rate $\delta = 0.1$ from

$$\int_{\mathcal{C}_{th}(\alpha)}^{\infty} dx F(x; 4M'_{i,max}, 0) = \alpha/n , \quad (25)$$

$$\int_0^{\mathcal{C}_{th}(\alpha)} dx F[x; 4M'_{i,max}, \lambda_{th}(\alpha, \delta)] = \delta , \quad (26)$$

where $F(x; k, \lambda)$ is the probability density function of the $\chi_k^2(\lambda)$ distribution and $4M'_{i,max}$ is the maximum number of sidebands for all f'_i in the window.

To simulate the gravitational wave search, we choose ρ_0^2 and draw $2\mathcal{F}$ values from the $\chi_4^2[\rho^2(f'_i)]$ distribution for each frequency bin f'_i . The search windows are constructed to ensure the entire sideband structure lies completely within a single window. Sidebands are also present in adjacent windows, however we focus only on a single window, assuming (conservatively) that all other windows contain only noise. Equation (24) makes the same assumption when computing the detection thresholds against which $HC_{2\mathcal{F}}$ and $HC_{\mathcal{C}}$ are compared. We conduct a 100 Monte-Carlo simulations for each ρ_0^2 and find the fraction where the computed statistic lies above its threshold value, which we call the detection rate.

Figure 5 displays detection rates for \mathcal{C} , $HC_{2\mathcal{F}}$, and $HC_{\mathcal{C}}$ as functions of increasing wave strain h/h_{th} , where h_{th} is defined as the wave strain corresponding to λ_{th} . For $\delta = 0.1$ (detection rate of 90%), $HC_{\mathcal{C}}$ detects wave strains 6% smaller than \mathcal{C} alone and $HC_{2\mathcal{F}}$ detects wave strains ~ 7

times greater than \mathcal{C} . Additionally, higher criticism on $2\mathcal{F}$ compared to \mathcal{F} provides a advantage comparable to $HC_{\mathcal{C}}$ over \mathcal{C} (not displayed in Figure 5). The results in Figure 5 assume perfect knowledge of the binary orbital period P . In summary, therefore, second-pass higher criticism on \mathcal{C} slightly boosts the sensitivity of the \mathcal{C} -statistic, while second-pass higher criticism on $2\mathcal{F}$ is not competitive. Both higher criticism statistics, in particular $HC_{2\mathcal{F}}$, are more robust than \mathcal{C} , as demonstrated in the next section.

4.3. Robustness of higher criticism

One significant advantage of higher criticism is its robustness when applied to signals that differ from the assumed form, e.g. due to phase wandering. In this section, we explore the performance of the statistics \mathcal{C} , $HC_{\mathcal{C}}$, and $HC_{2\mathcal{F}}$ when the form of the signal remains unaltered but the true signal parameters differ from those assumed in the search. For a comb search targeting a binary source, the orbital period P determines the sideband spacing of the \mathcal{F} -statistic in the frequency domain. An error ΔP between the assumed and true value of P produces a comb that does not coincide with the actual sidebands. The \mathcal{C} -statistic sums noise at each sideband whose frequency does not match a true sideband location, lowering the sensitivity. In contrast, $HC_{2\mathcal{F}}$ only assumes that all sidebands are located within a particular frequency window. It does not rely on knowledge of the precise location of each sideband and consequently is unaffected by an error in P . Both $HC_{\mathcal{C}}$ and the \mathcal{C} -statistic lose sensitivity as ΔP increases, however one would expect the relative advantage of $HC_{\mathcal{C}}$ compared to \mathcal{C} is expected to remain.

Observations of the orbital period of Sco X-1 have reported conflicting values for P . Gottlieb et al. (1975) measured $P = 68023.8 \pm 0.09$ s from archival optical observations. Recently, Vanderlinde et al. (2003) measured $P = 68163.6 \pm 8.6$ s with the Rossi X-ray Timing Explorer, but did not observe any significant periodicity at or near the previously reported value of 68023.8 s. Other LMXBs have larger uncertainty in their orbital parameters [e.g., see Tables 2–4 in Watts et al. (2008)]. In light of the potential uncertainty in P , the robustness of higher criticism is valuable.

We repeat the simulations described in Section 4.2 with the addition of an error ΔP between the assumed value of P used to construct the comb $q(f)$, and the true value used to generate the noncentrality parameters $\rho^2(f)$ for the \mathcal{F} -statistic at every sideband. Figure 6 displays detection rates with \mathcal{C} , $HC_{2\mathcal{F}}$, and $HC_{\mathcal{C}}$ as ΔP increases. As expected, $HC_{2\mathcal{F}}$ is unaffected by ΔP because higher criticism does not use any information about the sideband locations. \mathcal{C} and $HC_{\mathcal{C}}$ both lose sensitivity as ΔP increases, but the relative advantage of $HC_{\mathcal{C}}$ over \mathcal{C} increases from 6% when $\Delta P = 0$ to 15% for $\Delta P = 1$ s, 21% for $\Delta P = 3$ s, and 29% for $\Delta P = 9$ s. The difference between \mathcal{C} and $HC_{2\mathcal{F}}$ decreases from a factor of 7 when $\Delta P = 0$ to 5 for $\Delta P = 1$ s, 3 for $\Delta P = 3$ s, and 2 for $\Delta P = 9$ s. When $\Delta P = 0$, the \mathcal{C} -statistic peaks at the signal frequency f_0 . For $\Delta P > 0$, the error between the assumed and actual location of the sidebands accumulates to the point where the outermost sidebands no longer coincide with the assumed comb. This reduces the total number of signal sidebands summed by the \mathcal{C} -statistic, reducing its maximum amplitude (at




Fig. 5.— Monte-Carlo detection rate ($1-\delta$), equal to the fraction of simulations where the detection statistic lies above the detection threshold, as a function of wave strain, for the \mathcal{C} -statistic (solid black), $HC_{\mathcal{C}}$ (dashed red), and $HC_{2\mathcal{F}}$ (dotted blue).




Fig. 6.— Monte-Carlo detection rate $(1-\delta)$, equal to the fraction of simulations where the detection statistic lies above detection threshold, as a function of wave strain for the \mathcal{C} -statistic (solid black), $HC_{\mathcal{C}}$ (dashed red), and $HC_{2\mathcal{F}}$ (dotted blue). The error ΔP between assumed and true binary orbital periods increases from top to bottom: $\Delta P = 0, 1, 3, 9$ s.

f_0). Additionally, the \mathcal{C} -statistic peak broadens compared with $\Delta P = 0$; this maximum amplitude requires less signal bins summed and can now be reached at a number of the central sidebands. This situation moves closer to multiple, equal-strength signals, the problem for which higher criticism is designed, explaining the increasing advantage of $HC_{\mathcal{C}}$ over \mathcal{C} . In contrast, the \mathcal{C} -statistic requires a single source above threshold for detection. Interestingly, the sensitivity decrease of \mathcal{C} and $HC_{\mathcal{C}}$ stalls for $9 \text{ s} \lesssim \Delta P \lesssim 1000 \text{ s}$, when the accumulated error at the outermost sidebands is at least the sideband spacing. At this point, sidebands begin to overlap with other parts of the comb, shifted by integer multiples of the sideband spacing. As ΔP increases beyond 9 s, the number of sidebands correctly located by the template continues to decrease but is compensated for by a corresponding increase in the number of integer overlap sidebands.

5. Phase wandering

An accreting neutron star whose phase wanders in response to a variable accretion torque emits gravitational wave power in many frequency bins. Resampling is difficult as the phase model is usually unknown, e.g. there may be an offset between the radio/X-ray ephemeris and gravitational wave signal. Higher criticism can handle phase wandering robustly.

In its simplest form, a semi-coherent \mathcal{F} -statistic combines N coherent time intervals of equal length according to,

$$\mathcal{F}_{sc}(f) = \sum_{i=1}^N \mathcal{F}_i(f) , \quad (27)$$

where \mathcal{F}_i is the usual \mathcal{F} -statistic for the i -th interval. This method assumes a monochromatic source or, with some modification, a source whose phase evolves in a known way. Abbott et al. (2008a) described three semi-coherent methods to search for periodic gravitational waves in LIGO data. The methods cannot be applied to a source whose phase wanders unpredictably, e.g., due to accretion torque. We apply higher criticism to reanalyze detection statistics from multiple time intervals with this situation in mind.

We define n_{bins} to be the number of frequency bins and n_{GW} to be the number of bins containing a gravitational wave signal, assuming for simplicity that all signal bins have equal amplitude. In practical terms, this situation corresponds to n_{GW} distinct sources of similar strength, or n_{GW} orbitally modulated sidebands of a source in a binary system. We also assume that n_{bins} and n_{GW} are the same in each time interval, therefore $\epsilon = n_{\text{GW}}/n_{\text{bins}}$ is constant. We allow for the signal to jump arbitrarily between bins (not necessarily adjacent ones) from one time interval to the next, e.g., due to phase wandering.

For each interval, we obtain a $2\mathcal{F}$ value for each of the n_{bins} frequency bins. The n_{GW} signal bins follow a $\chi_4^2(\rho^2)$ distribution, the remainder obey a $\chi_4^2(0)$ distribution. Combining all $2\mathcal{F}$ values for the N intervals gives a total of $n = Nn_{\text{bins}}$ values. Converting from ϵ and n to r and β through equations (5) and (9), we use the theoretical detection boundary [equation (6)] to find

$\rho_{HC}^2(N, n_{GW}, n_{bins})$, the minimum noncentrality detectable with HC over N time intervals, if n_{GW} of the n_{bins} bins contain signal. This can be converted to a wave strain using equation (12). The \mathcal{F} -statistic threshold noncentrality $\rho_{\mathcal{F}}^2$ for a single interval is calculated similar to the \mathcal{C} -statistic [equations (25) and (26)] (Wette 2012),

$$\int_{\mathcal{F}_{th}}^{\infty} dx F(x; 4, 0) = \alpha/n , \quad (28)$$

$$\int_0^{\mathcal{F}_{th}} dx F(x; 4, \rho_{\mathcal{F}}^2) = \delta^{1/n_{GW}} . \quad (29)$$

The semi-coherent \mathcal{F} -statistic threshold $\rho_{\mathcal{F}_{sc}}^2$ can also be calculated from equations (28) and (29), substituting $4N$, rather than 4, degrees of freedom and assuming a monochromatic source. The semi-coherent \mathcal{F} -statistic cannot be applied to source whose phase wanders unpredictably and the threshold reverts to $\rho_{\mathcal{F}}^2$.

Comparing detection thresholds for HC and \mathcal{F} is difficult because the detection boundary for HC , derived in the limit $n \rightarrow \infty$, is not equivalent to the \mathcal{F} -statistic thresholds, computed for a specific false alarm and false dismissal probability. Furthermore, the HC detection boundary underestimates the threshold for finite n (see Section 2.4). We therefore recalibrate ρ_{HC}^2 by a constant such that $\rho_{HC}^2(N = 1, n_{GW} = 1, n_{bins}) = \rho_{\mathcal{F}}^2$, thereby arranging that both HC and \mathcal{F} have equal power when $N = n_{GW} = 1$. To test this recalibration, we conduct Monte-Carlo simulations to determine ρ_{HC}^2 for $\alpha = 0.01$, $\delta = 0.1$ when $N = n_{GW} = 1$ for $n_{bins} = 10^3, 10^4, 10^5, 10^6$. We find that the Monte-Carlo ρ_{HC}^2 agrees with $\rho_{\mathcal{F}}^2$ computed from equations (28) and (29) to within 1%. However, it is important to mind the difference between each threshold when interpreting the following results.

Figures 7 and 8 display ρ^2 required for detection with HC (thick curves) and \mathcal{F} (thin solid and dashed curves) for different combinations of N , n_{GW} , n_{bins} , and $\alpha = 0.01$, $\delta = 0.01$. Figure 7 displays the thresholds ρ_{HC}^2 and $\rho_{\mathcal{F}}^2$ as functions of n_{GW} for $(n_{bins}, N) = (10^6, 1), (10^6, 20), (10^9, 1), (10^9, 20)$. Figure 8 displays ρ^2 as a function of N for $(n_{bins}, n_{GW}) = (10^6, 1), (10^6, 10), (10^9, 1), (10^9, 10^3)$. For example, the search described in Section 4 has $n_{bins} \sim 10^9$ and $n_{GW} \sim 10^4$ (the number of orbital sidebands). The semi-coherent \mathcal{F}_{sc} is, in general, the most sensitive but as described above we are most interested in sources with unknown phase wandering, for which semi-coherent searches are not suitable. In all cases shown, HC outperforms \mathcal{F} . In Figure 7, we observe reduction in ρ_{HC}^2 by a factor of 1.5 to 2.9 for $n_{GW} = 10$ and 2.3 to 7.0 for $n_{GW} = 100$, compared to ρ_{HC}^2 when $n_{GW} = 1$. In Figure 8, we observe reduction in ρ_{HC}^2 by a factor of 1.4 to 2.2 for $N = 10$ and between 2.1 and 3.0 for $N = 100$, compared to $N = 1$.

The robustness of HC described in Section 5 allows one to combine data from multiple time intervals to increased the sensitivity, even for a source whose phase wanders unpredictably. In principle, one might improve detectability while remaining robust by including some simple information about the source frequency evolution, e.g. limiting wandering to some physically motivated range. In Figures 7 and 8, we assume no correlation, allowing the unlikely possibility that the




Fig. 7.— Noncentrality thresholds ρ_{HC}^2 (thick curves) and $\rho_{\mathcal{F}}^2$ (thin curves) as functions of n_{GW} for $(n_{\text{bins}}, N) = (10^6, 1), (10^6, 20), (10^9, 1), (10^9, 20)$




Fig. 8.— Noncentrality thresholds ρ_{HC}^2 (thick curves) and $\rho_{\mathcal{F}}^2$ (thin curves) as functions of N for $(n_{\text{bins}}, n_{\text{GW}}) = (10^6, 1), (10^6, 10), (10^9, 1), (10^9, 10^3)$

frequency wanders across the entire range covered by the search. One simple method is to divide the frequency range into intervals wide enough to contain the source throughout the observation and calculate HC for each interval.

6. Discussion

Higher criticism is a recently formulated statistical method designed to detect the presence of a sparse collection of signals too weak to be detected individually. In this paper, we apply higher criticism as a second pass to reanalyze gravitational wave search statistics and explore the feasibility of applications in two contexts: a targeted binary search (e.g. LMXB) and a phase-wandering source (e.g., glitching pulsar).

One advantage of higher criticism is its robust nature; it accommodates deviation from the expected phase evolution that hamper other search statistics. However, there is a trade-off. Higher criticism neglects some of the additional information used by other search methods which increase their performance under ideal conditions when the phase model is known.

In Section 4, we compare the performance of higher criticism and the \mathcal{C} -statistic for a targeted binary search. The \mathcal{C} -statistic is more sensitive than higher criticism applied to \mathcal{F} -statistic values by a factor ~ 7 , however, higher criticism applied to \mathcal{C} -statistic values gives a second-pass improvement of 6% over the \mathcal{C} -statistic. Furthermore, higher criticism is more robust to an error ΔP between the true and assumed (from observation) binary orbital period. While the absolute sensitivity of both \mathcal{C} and $HC_{\mathcal{C}}$ decrease as ΔP increases, the performance of $HC_{\mathcal{C}}$ relative to \mathcal{C} increases to a sensitivity improvement of 16% for $\Delta P = 1$ s and 29% for $\Delta P = 9$ s.

In Section 5, we consider a phase wandering source. The robustness of HC allows one to combine data from multiple time intervals and boost the sensitivity compared to a single interval, even for a unpredictable source. The noncentrality threshold decreases by a factor $\gtrsim 1.4$ over 10 intervals and a factor $\gtrsim 2$ over 100 intervals.

Another candidate for higher criticism is an all-sky search for unknown periodic sources. (We tested this in Appendix A). The results indicate that higher criticism provides no advantage over the \mathcal{F} -statistic, the statistics perform similarly. This stems from the distribution of source amplitudes: if neutron stars are uniformly distributed across the Galaxy, the closest, strongest source dominates the detectability. There is no advantage using higher criticism, which is designed to detect a group of signals, when there is effectively only one signal present.

While there is reason to be optimistic about possible applications of higher criticism, we draw attention to a number of significant concerns. Chief among these, higher criticism relies on knowing the background noise distribution. In this paper we adopted the common, simplifying assumption of stationary, Gaussian detector noise. In reality, detector noise is more complicated. However, it is well studied and modified forms of higher criticism have been developed for correlated noise (Hall

& Jin 2010) or when the noise distribution is imperfectly known (Delaigle et al. 2011).

Additionally, higher criticism only detects the presence of a group of signals. It cannot directly identify an individual source, which remains the ultimate goal of gravitational wave detection. However, given the cheap computational cost of applying higher criticism as a second pass to reanalyze already computed search statistics, and the advantages related to its robustness, it has the potential to complement and enhance existing and future searches, especially as a guide to where to look harder in parameter space for just-too-weak sources.

We thank L. Sammut, V. Dergachev, and members of the Continuous Wave Search Group of the LIGO Scientific Collaboration for helpful discussions. MFB was supported by an Australian Postgraduate Award. AM, AD and PH were supported by grants and fellowships from the Australian Research Council.

REFERENCES

- Abadie, J., Abbott, B. P., Abbott, R., et al. 2010a, *ApJ*, 722, 1504
- . 2010b, *Classical and Quantum Gravity*, 27, 173001
- . 2011, *ApJ*, 737, 93
- Abbott, B., Abbott, R., Adhikari, R., et al. 2004, *Phys. Rev. D*, 69, 082004
- . 2005, *Phys. Rev. D*, 72, 102004
- . 2007a, *Phys. Rev. D*, 76, 082001
- . 2007b, *Phys. Rev. D*, 76, 042001
- . 2008a, *Phys. Rev. D*, 77, 022001
- . 2008b, *ApJ*, 683, L45
- . 2009a, *Phys. Rev. D*, 79, 022001
- Abbott, B. P., Abbott, R., Adhikari, R., et al. 2009b, *Physical Review Letters*, 102, 111102
- . 2009c, *Phys. Rev. D*, 80, 042003
- . 2009d, *Reports on Progress in Physics*, 72, 076901
- . 2012, *Journal of Instrumentation*, 7, 3012
- Arnett, W. D., Schramm, D. N., & Truran, J. W. 1989, *ApJ*, 339, L25

- Bildsten, L. 1998, *ApJ*, 501, L89+
- Brady, P. R., & Creighton, T. 2000, *Phys. Rev. D*, 61, 082001
- Cayón, L., Jin, J., & Treaster, A. 2005, *MNRAS*, 362, 826
- Chakrabarty, D., Morgan, E. H., Muno, M. P., et al. 2003, *Nature*, 424, 42
- Cook, G. B., Shapiro, S. L., & Teukolsky, S. A. 1994, *ApJ*, 424, 823
- Cutler, C. 2011, arXiv:1104.2938
- Cutler, C., Gholami, I., & Krishnan, B. 2005, *Phys. Rev. D*, 72, 042004
- Cutler, C., & Schutz, B. F. 2005, *Phys. Rev. D*, 72, 063006
- Delaigle, A., & Hall, P. 2009, in *Perspectives in mathematical sciences I: Probability and Statistics*, ed. B. R. N. Sastry, M. Delampady & T. Rao (World Scientific Publishing), 109–138
- Delaigle, A., Hall, P., & Jin, J. 2011, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73, 283
- Donoho, D., & Jin, J. 2004, *Annals of Statistics*, 32, 962
- Gottlieb, E. W., Wright, E. L., & Liller, W. 1975, *ApJ*, 195, L33
- Hall, P., & Jin, J. 2010, *Annals of Statistics*, 38, 1686
- Jaranowski, P., Królak, A., & Schutz, B. F. 1998, *Phys. Rev. D*, 58, 063001
- Krishnan, B., Sintes, A. M., Papa, M. A., et al. 2004, *Phys. Rev. D*, 70, 082001
- Manchester, R. N., Hobbs, G. B., Teoh, A., & Hobbs, M. 2005, *AJ*, 129, 1993
- Mastrano, A., & Melatos, A. 2012, *MNRAS*, 421, 760
- Melatos, A., & Payne, D. J. B. 2005, *ApJ*, 623, 1044
- Messenger, C., & Woan, G. 2007, *Classical and Quantum Gravity*, 24, 469
- Ushomirsky, G., Cutler, C., & Bildsten, L. 2000, *MNRAS*, 319, 902
- Vanderlinde, K. W., Levine, A. M., & Rappaport, S. A. 2003, *PASP*, 115, 739
- Vigelius, M., & Melatos, A. 2009, *MNRAS*, 395, 1972
- Watts, A. L., Krishnan, B., Bildsten, L., & Schutz, B. F. 2008, *MNRAS*, 389, 839
- Wette, K. 2012, *Phys. Rev. D*, 85, 042003

A. All-sky search

One might ask whether higher criticism can be applied profitably to all-sky searches. After all, from an estimated Galactic population of 10^9 neutron stars (Arnett et al. 1989), only ~ 2000 pulsars have been discovered as radio sources, of which $\sim 10\%$ have frequencies in the range where current interferometers are most sensitive, $f_0 > 100$ Hz (Manchester et al. 2005)¹. Theories of neutron star quadrupoles suggest that only a handful of the observed pulsars have any chance of being detectable with current sensitivity limits (Mastrano & Melatos 2012), but many of the radio-quiet objects are much closer to the Earth than their radio-loud brethren. All-sky searches cover a range of frequencies at many sky locations, hoping to detect one of these unknown sources. A number of all-sky searches have been carried out (Abbott et al. 2005, 2007a, 2008a, 2009a,b,c), but no detection has been announced. The searches typically use the \mathcal{F} -statistic, or related statistics. Hence they are candidates for reanalysis with higher criticism.

In contrast to the other applications in this paper, the sources targeted by an all-sky search have a power law (non-uniform) wave strain distribution. We follow Cutler & Schutz (2005) and assume neutron stars are spread uniformly throughout the Galaxy. Then the distance distribution separates into two parts: a local uniform three-dimensional distribution up to the thickness of the Galactic disk, and a uniform two-dimensional distribution beyond. Taking the disk to be 600 pc thick and 10 kpc in diameter, the 3D and 2D distributions correspond to $D < 300$ pc and 300 pc $< D < 5$ kpc respectively, where D is the distance from Earth. We further assume that all sources have the same intrinsic amplitude, so that h_0 at Earth depends on distance alone. Changing variables from distance D to noncentrality parameter ρ^2 , [see equation (10); also Cutler & Schutz (2005) for details] we arrive at the distribution,

$$\sigma(\rho^2) = \begin{cases} n_{3D}[\rho^2]^{-5/2}, & (3D) , \\ n_{2D}[\rho^2]^{-2}, & (2D) , \end{cases} \quad (A1)$$

where n_{2D} and n_{3D} are normalization constants which depend on h_0 , $S_h(f)$ and T_{obs} through equation (10).

To illustrate, we consider the same search parameters as a completed LIGO all-sky search for periodic sources. Abbott et al. (2007a) searched the 10 hr of LIGO data with the best sensitivity from the second science run and computed the \mathcal{F} -statistic for 1.6×10^8 frequency bins at each of 3×10^4 sky locations, a total of $n = 5 \times 10^{12}$ \mathcal{F} -statistic values. We estimate the number of neutron stars with frequencies that fall in the range of the search (160-728.8 Hz) to be 10^7 (of an estimated Galactic population of 10^9). Under these assumptions, each of the 5×10^{12} frequency bins is equally likely to contain some signal, with $\epsilon = 2 \times 10^{-6}$.

We compute the Monte-Carlo higher criticism detection rate as a function of $\rho_{max}^2/\rho_{\mathcal{F}}^2$, where $\rho_{\mathcal{F}}^2$ is the noncentrality threshold for a \mathcal{F} -statistic detection with 1% false alarm rate and 10%

¹<http://www.atnf.csiro.au/research/pulsar/psrcat>

false dismissal rate [see equations (28) and (29)]. We can then compare the noncentrality of the brightest source, ρ_{max}^2 , required for detection with higher criticism with the equivalent quantity for a \mathcal{F} -statistic search. To save computation, we scale simulations from $n = 5 \times 10^{12}$ to $n' = 10^6$ by keeping parameters r and β constant in equations (5) and (9) and converting from ϵ and ρ^2 for $n = 5 \times 10^{12}$ to ϵ' and $\rho^{2'}$ for $n' = 10^6$.

Signal values are drawn from a $\chi_4^2(\rho^{2'})$ distribution with noncentrality parameter $\rho^{2'}$ generated according to equation (A1). To fix n_{2D} and n_{3D} we specify ρ_{max}^2 for a source at a distance of 10 pc. The remaining \mathcal{F} values, containing only noise, are drawn from the $\chi_4^2(0)$ distribution, and we compute HC for the \mathcal{F} -statistic values. For each ρ^2 , we run 100 simulations and the detection rate equals the fraction of simulations with HC or \mathcal{F} above their corresponding thresholds for detection. Finally, we convert noncentrality ρ^2 to wave strain h through equation (12) to plot a more meaningful quantity.

Figure 9 displays the Monte-Carlo detection rate as a function of $h/h_{\mathcal{F}}$, where h is the wave strain for a source at 10 pc, corresponding to ρ_{max}^2 , normalized by $h_{\mathcal{F}}$, the wave strain required for detection with the \mathcal{F} -statistic at 10 pc. Figure 9 shows that higher criticism provides no advantage over the \mathcal{F} -statistic; they perform almost identically. The fluctuations in Figure 9 relate to the randomly sampled source distribution. Despite the large number of sources, $\sigma(\rho^2)$ in equation (A1) falls off too fast with ρ^2 and effectively only the brightest source, rather than the collection of sources, is detected. To test this we repeated the above simulations using only the single strongest source (replacing the rest with noise) and found the same result. The similar performance of \mathcal{F} and $HC_{\mathcal{F}}$ agrees with Section 5, where Monte-Carlo simulations found HC and \mathcal{F} had equal thresholds when only one source was present. Therefore, for an all sky search, higher criticism does not provide any benefit over a first pass \mathcal{F} -statistic search.




Fig. 9.— Monte-Carlo detection rate with higher criticism (thick dashed red) versus \mathcal{F} -statistic (thin solid black) as a function of $h/h_{\mathcal{F}}$, where h is the wave strain of a source located at 10 pc and $h_{\mathcal{F}}$ is the strain required by a source at 10 pc for detection with the \mathcal{F} -statistic.