

Fully coherent follow-up of continuous gravitational-wave candidates: An application to Einstein@Home results

M. Shaltev,¹ P. Leaci,² M. A. Papa,^{2,3} and R. Prix¹¹*Albert Einstein Institut, Callinstraße 38, 30167 Hannover, Germany*²*Albert Einstein Institut, Am Mühlenberg 1, 14476 Potsdam-Golm, Germany*³*University of Wisconsin Milwaukee, 1900 E. Kenwood Boulevard, 23201 Milwaukee, USA*

(Received 8 May 2014; published 25 June 2014)

We characterize and present the details of the follow-up method used on the most significant outliers of the Hough Einstein@Home all-sky search for continuous gravitational waves [J. Aasi *et al* Phys. Rev. D **87**, 042001 (2013)]. This follow-up method is based on the two-stage approach introduced by [M. Shaltev and R. Prix, Phys. Rev. D **87**, 084057 (2013)], consisting of a semicoherent refinement followed by a fully coherent zoom. We quantify the efficiency of the follow-up pipeline using simulated signals in Gaussian noise. This pipeline does not search beyond first-order frequency spin-down, and therefore we also evaluate its robustness against second-order spin-down. We present the details of the Hough Einstein@Home follow-up [J. Aasi *et al* Phys. Rev. D **87**, 042001 (2013)] on three hardware-injected signals and on the eight most significant search outliers of unknown origin.

DOI: 10.1103/PhysRevD.89.124030

PACS numbers: 04.30.-w

I. INTRODUCTION

The search for unknown sources of continuous gravitational waves (CWs) is computationally bound due to the enormous parameter space that needs to be covered [1]. Advanced semicoherent search techniques, such as [2,3], are typically used to identify interesting regions of the parameter space, which then require fully coherent follow-up studies in order to confirm or discard potential CW candidates. The parameter space associated with these candidates is substantially smaller than the original search space. However, it is still large enough to lead to a prohibitive computing cost, when data of order of months or years are analyzed fully coherently with a classical grid-based method [4]. Therefore, an alternative follow-up method was developed which combines the \mathcal{F} -statistic [5,6] with a mesh adaptive direct search (MADS) [7] algorithm. This allows us to fully coherently examine long data sets at a feasible computational cost [8].

In the present work we describe how the two-stage algorithm proposed in [8] was adapted to follow up on the most significant outliers in the Hough S5 Einstein@Home search [9]. We first validate the follow-up pipeline by performing Monte Carlo studies. We inject and search for simulated CW signals added into simulated Gaussian noise data. Then we show how the search method was applied to 35 outliers identified in the Einstein@Home search, 27 of which are associated with 3 simulated signals (hardware injections, as discussed in Sec. VA).

The paper is organized as follows. In Sec. II we briefly recap the Einstein@Home all-sky search for periodic gravitational waves in data from the fifth Laser Interferometer Gravitational-wave Observatory (LIGO) science run (S5). In Sec. III we summarize the two-stage follow-up method and

introduce the search pipeline. The efficiency of the follow-up algorithm is tested with Monte Carlo studies presented in Sec. IV. In Sec. VA we present the follow-up results for the 27 outliers associated with 3 hardware injections. In Sec. VB we show the results of the follow-up for the remaining CW outliers. Section VI presents a discussion of the results and concluding remarks.

A. Notation and conventions

When referring to a quantity Q of the original Hough search, we denote it as Q_{HS} . A quantity measured after the prerefinement stage is denoted as Q_{PR} , after the refinement as Q_{R} , and as \tilde{Q}_{Z} , after the fully coherent zoom stage using all the data (consistent with the notation of [8]). We use an overbar (\bar{Q}) to denote an average over segments.

II. THE HOUGH S5 EINSTEIN@HOME ALL-SKY SEARCH

The Einstein@Home all-sky search [9] uses the semicoherent Hough-transform method [10], which consists of dividing the entire data span into N shorter segments of duration ΔT . In a first step, a coherent \mathcal{F} -statistic search is performed on a coarse grid for each of the data segments. Then the Hough number-count statistic, defined in Eq. (6), is computed on a finer grid, using the \mathcal{F} -statistic values from the individual segments.

In this paper we focus on the S5R5 search of [9], which spans approximately 264 days of data from the Hanford (H1) and Livingston (L1) LIGO detectors. This dataset was divided into $N = 121$ segments of duration $\Delta T = 25$ hours each. The parameter space covered by this search spans the entire sky, a frequency range $f \in [50, 1190]$ Hz, and a spin-down range $\dot{f} \in [-20, 1.1] \times 10^{-10}$ Hzs⁻¹.

The phase evolution of the expected signal at the detector can be written as [5]

$$\Phi(t) \approx \Phi_0 + 2\pi \sum_{k=0}^s \frac{f^{(k)}(t_0)(t-t_0)^{k+1}}{(k+1)!} + 2\pi \frac{\mathbf{r}(t)}{c} \cdot \mathbf{n} \sum_{k=0}^s \frac{f^{(k)}(t_0)(t-t_0)^k}{k!}, \quad (1)$$

where Φ_0 is the initial phase, $f^{(k)} \equiv \frac{d^k f}{dt^k}$ represents the time derivatives of the signal frequency f at the solar system barycenter (SSB) at reference time t_0 , s is the maximal considered spin-down order, c is the speed of light, and $\mathbf{r}(t)$ is the vector pointing from the SSB to the detector. The unit vector $\mathbf{n} \equiv (\cos \alpha \cos \delta, \sin \alpha \cos \delta, \sin \delta)$ points from the SSB to the CW source, where α and δ are the standard equatorial coordinates referring to right ascension and declination, respectively.

The \mathcal{F} -statistic is one of the standard coherent techniques used to extract the CW signals from the noisy detector data. This statistic is the result of matched filtering the data with a signal template characterized by the phase-evolution parameters $\lambda \equiv \{\alpha, \delta, f, \dot{f}\}$. The amplitude parameters, namely the intrinsic amplitude h_0 , the inclination angle ι , the polarization angle ψ , and the initial phase ϕ_0 , have been analytically maximized over [5]. In a coherent grid-based \mathcal{F} -statistic search, the number of templates increases with a high power of the observation time [11]. Hence these searches are not suitable for wide parameter-space all-sky surveys. However, the reduction of the coherent baseline in a semicoherent search [1,2] makes these techniques computationally feasible in a distributed computing environment such as Einstein@Home and (usually) more sensitive at fixed computing cost [12].

The template bank used to cover the parameter space is constructed using the notion of mismatch [13,14]. This is defined as the fractional loss of squared signal-to-noise ratio (SNR) between a template λ and the signal location λ_s . We use the definition of SNR given in [5] and denote it as ρ .

To quadratic order in parameter-space offsets $\Delta\lambda^i \equiv \lambda^i - \lambda_s^i$, the mismatch can be approximated by

$$\mu^* \equiv g_{ij}(\lambda_s) \Delta\lambda^i \Delta\lambda^j, \quad (2)$$

where g_{ij} is a symmetric positive-definite matrix referred to as the parameter-space metric. The indices i and j label the phase-evolution parameters, and we use summation convention over repeated indices. This metric mismatch μ^* can be interpreted as a distance measure in parameter space.

In the S5R5 analysis the templates at frequency f were placed on a coarse grid constructed using the following spacings [9]:

$$d\theta_{\mathcal{F}} = \frac{\sqrt{3}c}{v_d f \Delta T}, \quad df = \frac{\sqrt{12m}}{\pi \Delta T}, \quad \dot{f} = \frac{\sqrt{3.3m}}{\Delta T^2}, \quad (3)$$

where $d\theta_{\mathcal{F}}$ is the angular resolution of the coarse sky grid, df and \dot{f} are the frequency and spin-down resolutions, respectively; m is the nominal single-dimension mismatch, taken equal to 0.3 in [9]; and v_d is the Earth's rotation speed at the equator. Due to limitations of the Einstein@Home environment, in particular the amount of available random-access memory on the computing nodes, the spin-down resolution was not increased for the fine grid. Instead the \dot{f} resolution of Eq. (3) was determined in a Monte Carlo study so as not to significantly lose detection efficiency.

The resolution of the fine sky grid at frequency f is given by [9]

$$d\theta_H = \frac{cd\dot{f}}{\wp f v_y}, \quad (4)$$

where \wp is the pixel factor and v_y is the Earth's orbital velocity. With $\wp = 0.5$ and $m = 0.3$, the sky refinement used in the S5R5 search yields $\mathcal{N}_{\text{sky}}^{\text{ref}} = (d\theta_{\mathcal{F}}/d\theta_H)^2 \approx 8444$ [9].

Every parameter-space point of the search is assigned a significance, or critical ratio (CR), value [9]:

$$\text{CR} = \frac{n_c - \bar{n}_c}{\sigma}, \quad (5)$$

with

$$n_c = \sum_{\ell=1}^N w_{\ell} n_{\ell} \quad (6)$$

the Hough number count, where w_{ℓ} is the weight for segment ℓ at a frequency f and a sky position (α, δ) ; $n_{\ell} = 1$ if the \mathcal{F} -statistic crosses a certain threshold value (namely $2\mathcal{F} > 5.2$ in [9]), otherwise $n_{\ell} = 0$; and \bar{n}_c and σ are the expected value and the standard deviation of n_c in Gaussian noise. The candidates are ordered by their significance.

III. FOLLOW-UP METHOD

A. The modified two-stage follow-up

A slightly adapted version of the two-stage follow-up procedure [8] was used in [9] and is presented here. As mentioned in Sec. II, the original Hough search did not use refinement in f , and this led to a reduction in localization accuracy. To recover from this, we perform a *prerefinement* by rerunning the original Hough search with a finer grid around the outlier being followed up. Namely, we increase the resolution of the f grid by a factor $N = 121$,¹ and the sky resolution by doubling the pixel factor \wp in Eq. (4). The usefulness of this prerefinement is illustrated in Fig. 1.

The loudest parameter-space point after prerefinement provides the starting point for the subsequent MADS-based follow-up method described in [8]: namely, we first employ the semicoherent \mathcal{F} -statistic $2\mathcal{F}$, defined as

¹The refinement of the f -grid resolution is linear in the number of segments; see [15].

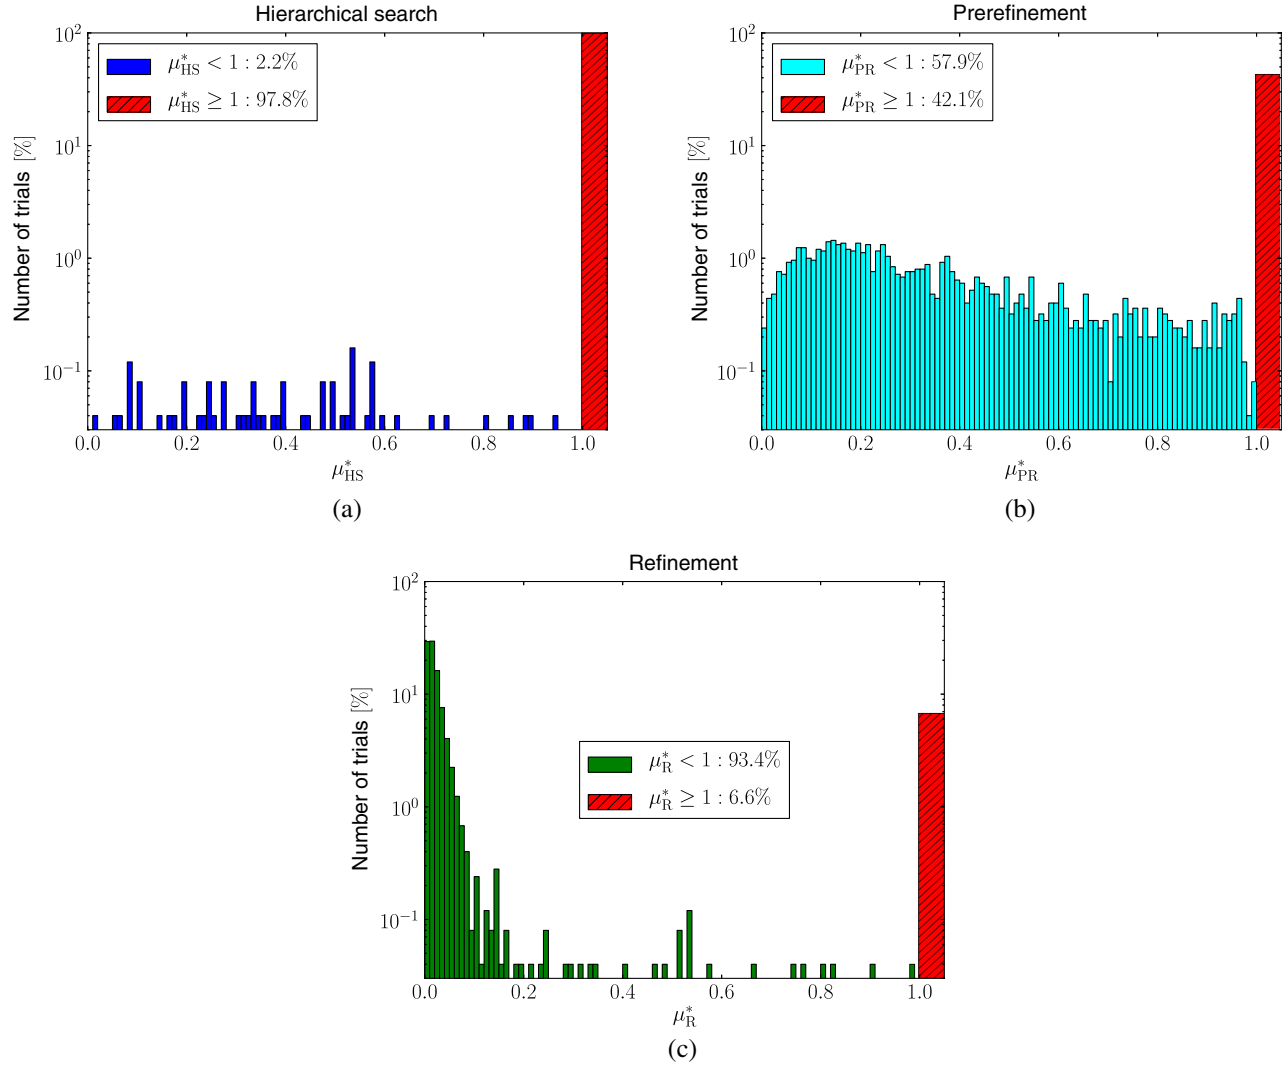


FIG. 1 (color online). Semicoherent metric mismatch of the subset of 2500 injections with $\overline{\rho^2} \in [5, 6]$ at different stages of the Monte Carlo. Panel (a) shows the metric mismatch distribution after the original hierarchical search. Panel (b) shows the metric mismatch distribution after the prerefinement stage. Panel (c) shows the metric mismatch distribution after the refinement stage. The hatched bar in each panel shows the percentage of trials with $\mu^* \geq 1$.

$$\overline{2\mathcal{F}}(\lambda) \equiv \frac{1}{N} \sum_{\ell=1}^N 2\mathcal{F}_{\ell}(\lambda), \quad (7)$$

where $2\mathcal{F}_{\ell}(\lambda)$ is the coherent \mathcal{F} -statistic computed on segment ℓ at the parameter-space point λ . This is computed on the original Hough segments to further improve the localization of the maximum-likelihood parameter-space point, using the gridless MADS search method described in more detail in [8]. This stage is called *refinement*, with detection statistic $\overline{2\mathcal{F}}_R$ for the loudest resulting candidate.

Next we apply the so-called \mathcal{F} -statistic consistency veto of [9,16], namely

$$\text{veto if } \overline{2\mathcal{F}}_R < \max\{\overline{2\mathcal{F}}_R^{\text{H1}}, \overline{2\mathcal{F}}_R^{\text{L1}}\}, \quad (8)$$

where $\overline{2\mathcal{F}}_R^{\text{H1,L1}}$ denote the corresponding semicoherent \mathcal{F} -statistic values from the individual detectors H1 and L1, respectively.

Then, in the so-called *zoom* stage, we compute the fully coherent $\widetilde{2\mathcal{F}}_Z$ statistic using all the data. From this we determine whether the resulting candidate is consistent with the signal model or with Gaussian noise.

B. Classification of zoom outcomes

We distinguish three possible outcomes of the zoom stage:

- (i) *Consistency with Gaussian noise (G)*—the fully coherent $\widetilde{2\mathcal{F}}_Z$ value does not exceed a prescribed threshold, i.e.,

$$\widetilde{2\mathcal{F}}_Z < \widetilde{2\mathcal{F}}_{\text{th}}^{(G)}, \quad (9)$$

where $\widetilde{2\mathcal{F}}_{\text{th}}^{(G)}$ is chosen to correspond to some (small) false-alarm probability p_{fA} in Gaussian noise.

The single trial false-alarm probability for a given $\widetilde{\mathcal{F}}_{\text{th}}^{(G)}$ threshold is $p_{\text{fA}}^1 = (1 + \mathcal{F})e^{-\mathcal{F}}$ (see, e.g., Ref. [8] for details). For example, we find that a threshold of $\widetilde{\mathcal{F}}_{\text{th}}^{(G)} = 90$ corresponds to a false-alarm probability $\sim \mathcal{O}(10^{-18})$ for a single template. Assuming \mathcal{N} independent templates and $p_{\text{fA}} \ll 1$, the false alarm is $p_{\text{fA}} \approx \mathcal{N}p_{\text{fA}}^1$.

- (ii) *Non-Gaussian origin* ($\neg G$)—the candidate is loud enough to be inconsistent with Gaussian noise at the chosen p_{fA} , i.e.,

$$\widetilde{\mathcal{F}}_Z \geq \widetilde{\mathcal{F}}_{\text{th}}^{(G)}. \quad (10)$$

- (iii) *Signal recovery* (S)—defined as a subclass of $\neg G$; namely, a signal is considered recovered if for the final zoomed candidate the $\widetilde{\mathcal{F}}_Z$ value exceeds the Gaussian-noise threshold $\widetilde{\mathcal{F}}_{\text{th}}^{(G)}$ and falls into a predicted signal interval:

$$\widetilde{\mathcal{F}}_{\text{th}}^{(S)} < \widetilde{\mathcal{F}}_Z < \widetilde{\mathcal{F}}_{\text{max}}^{(S)}, \quad (11)$$

where $\widetilde{\mathcal{F}}_{\text{th}}^{(S)} \equiv \max\{\widetilde{\mathcal{F}}_{\text{th}}^{(G)}, \widetilde{\mathcal{F}}_o - n_u \sigma_o\}$, and $\widetilde{\mathcal{F}}_{\text{max}}^{(S)} \equiv \widetilde{\mathcal{F}}_o + n_u \sigma_o$, with expectation

$$\widetilde{\mathcal{F}}_o \approx 4 + N(\overline{\mathcal{F}}_R - 4), \quad (12)$$

and variance

$$\sigma_o^2 \approx 2(4 + 2N(\overline{\mathcal{F}}_R - 4)). \quad (13)$$

The number n_u determines the probability that a true signal candidate would fall into this interval. For example, $n_u = 6$ corresponds roughly to a confidence of $\sim 99.6\%$ (provided G and S are disjoint).

C. Choice of MADS parameters

In both stages the parameter space is explored on a dynamically created mesh by using a MADS-based algorithm [7]. MADS itself is a general purpose algorithm for derivative-free optimization, which is typically applied to computationally expensive problems with unknown derivatives. The input to the MADS-based algorithm is a starting point λ_c , a search bounding box $\Delta\lambda_R$ around λ_c and a set of MADS parameters, which govern the choice of evaluation points. These parameters are $\{d\lambda, u_b, w_{\text{min}}^+, w_{\text{max}}^+, w^-, p\}$, where $d\lambda$ is the initial step, u_b is the mesh update basis, w_{min}^+ and w_{max}^+ are the mesh-coarsening exponents, w^- denotes the mesh-refining exponent² and p is the maximum number of templates to search over; for details, we refer the

²During the MADS optimization at iteration $k + 1$, the mesh spacing Δ is computed by $\Delta_{k+1} = u_b^{w^+} \Delta_k$, if a better point is found at iteration k , and by $\Delta_{k+1} = u_b^{w^-} \Delta_k$ otherwise [7].

TABLE I. Follow-up algorithm parameters for the refinement and zoom stage.

Stage	w^-	w_{min}^+	w_{max}^+	u_b	p
R	−1	1	20	2	20000
Z	−1	1	50	1.2	20000

reader to Sec. III E in [8]. The algorithm parameters for the MADS-based refinement and zoom stage are summarized in Table I. These parameters have been found to yield good results in Monte Carlo studies.

D. Follow-up parameter-space regions

We stress that the bounding box $\Delta\lambda_R$ used for the refinement differs with respect to what is described in [8]. There the refinement is restricted to the semicoherent metric ellipsoid centered on a candidate. Here instead, the refinement stage is performed on a box that was empirically determined, i.e., by performing Monte Carlo studies, to be large enough to contain the true signal location with very high confidence:

$$\begin{aligned} \Delta\alpha &= 0.4 \text{ rad}, & \Delta\delta &= 0.4 \text{ rad} \\ \Delta f &= 1 \times 10^{-4} \text{ Hz}, & \Delta\dot{f} &= 1 \times 10^{-9} \text{ Hz/s}. \end{aligned} \quad (14)$$

Given that this follow-up was not computationally limited, we did not attempt to find the smallest possible refinement region.

The zoom search is constrained by a Fisher ellipse scaled to 24 standard deviations, as described in [8]. This large number was chosen empirically by increasing it until the pipeline performance did not further improve.

The minimal spin-down order required is related to parameter-space thickness measured in terms of the extent of the metric ellipse along that direction [1,2,12]. As a rule of thumb, the maximal spin-down order required in a search increases with the time spanned by the data. In the Hough Einstein@Home all-sky search [9], the follow-up procedure did not include second-order spin-down. In Sec. IV we show the performance of the follow-up pipeline on signals with zero second-order spin-down, while in Sec. IV C we study the robustness of this method in the case of maximal second-order spin-down (as considered in [9]).

IV. MONTE CARLO STUDIES

A. Setup

We test the proposed follow-up pipeline in an end-to-end Monte Carlo study using the LALSuite [17] software package. In particular we use the following LALAPPS applications: MAKEFAKEDATA_v4 to generate Gaussian noise and inject CW signals; FSTATMETRIC_v2 to compute

the fully coherent or semicoherent metric; HIERARCHICAL SEARCH for the semicoherent Hough-transform search; FSTATSCNOMAD for the semicoherent \mathcal{F} -statistic optimization with MADS; and FSTATFCNOMAD for the fully coherent \mathcal{F} -statistic MADS optimization, where for the MADS algorithm we use the reference implementation NOMAD [18].

We apply the follow-up chain to 15000 different noise realizations with and without injected signals. The Gaussian noise realizations are generated with the MAKEFAKEDATA_V4 application using the same timestamps of the SFTs³ used in the original Einstein@Home search, with detector noise level of $\sim 2 \times 10^{-23}$ Hz^{-1/2} per detector. The signal parameters are uniformly drawn in the ranges $\cos i \in [-1, 1]$, $\psi \in [-\pi/4, \pi/4]$, $\phi_0 \in [0, 2\pi]$, and $f \in [185, 186]$ Hz, and the sky position is drawn isotropically on the sky. The frequency range has been chosen in the most sensitive region of the LIGO detectors. The spin-down value is randomly chosen in the range $\dot{f} \in (-\frac{f_0}{\tau_0}, 0.1 \frac{f_0}{\tau_0})$, with minimal spin-down age $\tau_0 = 800$ yr at $f_0 = 50$ Hz. The signal amplitude is high enough that the SNR² in the point of injection is uniformly distributed in the range $\bar{\rho}^2 \in [0, 6]$.

We begin the end-to-end validation with a simulation stage of the original S5R5 Einstein@Home search by using the original search setup, i.e., the same frequency and spin-down grid spacings given by Eq. (3). The S5R5 search has been partitioned in independent computing tasks, referred to as work units (WUs). For a detailed discussion of the WU, see Sec. III C in [9]. To save computing power, we do not rerun an entire WU in this simulation stage, but we center a search grid around a random point in the vicinity of the injected signal, searching over ten frequency bins in total. The sky grid is constructed by extracting 16 points around the candidate from the original sky-grid file. However, this reduced parameter-space size is still sufficiently large to make possible the selection of candidates due to the noise—if the signal is weak (as might happen in a real search)—and not artificially select a point close to the true signal location.

In Fig. 1 we show the semicoherent metric mismatch distribution, computed with Eq. (2), after the Hough search, the prerefinement Hough search, and the refinement stage. The loudest point selected from the refinement stage is used as a starting point for the fully coherent \mathcal{F} -statistic zoom search.

³SFT is the acronym used for short time baseline Fourier transform of the calibrated detector strain data. The duration of the SFTs is typically 1800 seconds. SFTs are used as inputs to many CW searches to construct the longer baseline of the coherent segments, e.g., the semicoherent Hough-transform search, as well as the fully coherent follow-up. See Ref. [19] for details on the implementation.

B. Efficiency of the follow-up pipeline

We first apply the follow-up pipeline to Gaussian noise data without any injected signals. This is required to ensure the applicability of the threshold $\widetilde{2\mathcal{F}}_{\text{th}}^{(G)} = 90$ used to consider a candidate to conform with the Gaussian noise hypothesis. The distribution of the $\widetilde{2\mathcal{F}}_Z$ values is plotted in Fig. 2. The maximal value found is $\widetilde{2\mathcal{F}}_Z^{\text{max}} = 79.36$, which is well below the $-G$ threshold of $\widetilde{2\mathcal{F}}_{\text{th}}^{(G)} = 90$.

In Fig. 3a we plot the percentage of the injected signals classified as recovered (S) and of non-Gaussian origin ($-G$), as a function of the average $\overline{2\mathcal{F}}$ value of the candidate after the simulation stage. We are able to distinguish $\geq 90\%$ of the candidates from Gaussian noise above $\overline{2\mathcal{F}}_c \gtrsim 6.0$, and we recover $\geq 90\%$ of the signals (S) for candidates with $\overline{2\mathcal{F}}_c \gtrsim 6.2$.

C. Robustness to second-order spin-down signals

The follow-up pipeline described in this work is limited to first-order spin-down in the signal model, which can lead to losses of SNR over long observation times for signals with nonzero second-order spin-down. In order to test the robustness of the follow-up method, we repeat the Monte Carlo simulation on signals with a fixed second-order spin-down value of $\dot{f} = 8 \times 10^{-20}$ Hz/s², which is estimated by using the minimal spin-down age τ_0 , namely f/τ_0^2 , as considered in [9]. The result of this simulation is presented in Fig. 3b and shows that for candidates with $\overline{2\mathcal{F}}_c \approx 6.5$, the follow-up pipeline is still able to distinguish close to 90% of the candidates from Gaussian noise. Given that this was the detection threshold used in the S5R5 search [9], we




FIG. 2 (color online). The figure shows the $\widetilde{2\mathcal{F}}_Z$ distribution after the fully coherent 4-D $\{\alpha, \delta, f, \dot{f}\}$ zoom stage of 15000 searches in pure Gaussian noise data without injected signals. The maximum value is $\widetilde{2\mathcal{F}}_Z^{\text{max}} = 79.36$, and the mean value is $\langle \widetilde{2\mathcal{F}}_Z \rangle = 51.40$ (dotted line). The labels 1,2,3 refer to the search outliers at roughly 434, 677, and 984 Hz, respectively (see Table III). The vertical red line marks the noise threshold.




FIG. 3 (color online). Monte Carlo study of the efficiency of the follow-up pipeline. Panel (a) shows the percentage of injected signals without second-order spin-down classified as recovered ($-S$), and with a non-Gaussian origin ($\times\neg G$), as function of the average $2\mathcal{F}$ value of the candidate after the original semicoherent Hough-transform search ($\overline{2\mathcal{F}_c}$). Panel (b) shows the percentage of injected signals with second-order spin-down classified as recovered ($-S$), and with a non-Gaussian origin ($\times\neg G$), as a function of $\overline{2\mathcal{F}_c}$. The error bars account for 95% confidence level. The 90% detection probability is marked with the horizontal dashed line. The vertical dotted line denotes the $2\mathcal{F}$ threshold used to select candidates in [9].

conclude that the follow-up of the resulting candidates did not substantially reduce the detection efficiency of the original search.

V. FOLLOW-UP OF S5R5 SEARCH OUTLIERS

In this section we report details on the follow-up of the S5R5 search outliers above the detection threshold of $\overline{2\mathcal{F}} \geq 6.5$, as originally reported in [9]. For practical purposes these search outliers were divided into two classes, depending on whether or not they are associated with hardware injections.

A. Search outliers associated with hardware injections

The CW hardware injections (referred to as “fake pulsars”) are simulated signals, physically added into the control system of the interferometer to produce a detector response similar to what should be generated if a CW is present. The aim of such injections is to test and validate analysis codes and search pipelines.

The S5R5 Einstein@Home search [9] identified three fake pulsars, referred to as Pulsars 2, 3, and 5. In this section we detail the follow-up of the search outliers associated with these hardware injections. Each injection typically produced many significant outliers. We apply a simple clustering algorithm in order to follow up on only the most interesting ones. Namely, for each hardware injection, we identify the loudest outlier and remove all neighboring search outliers falling into the refinement box given in Eq. (14). We repeat this procedure until there are no more search outliers left. A similar clustering algorithm was developed for the galactic-center search [20,21].

There are, for instance, 88 parameter-space points associated with Pulsar 2 injected at ~ 575 Hz. After the clustering procedure, the number of search outliers to follow up on is reduced to 16. For Pulsar 3, injected at ~ 108 Hz, the number of parameter-space points to follow up on shrinks from 80 to 9. For Pulsar 5, injected at ~ 52 Hz, there are only two search outliers, which fall into different search boxes and are therefore unaffected by the clustering.

In Table II we summarize, for each fake pulsar, the recovered parameters of the loudest outlier resulting from the follow-up. All of the injections were recovered at parameter-space points very close to the injected signal parameters, as quantified by the values of the metric mismatch μ_Z^* . We note that the recovered detection statistic $2\mathcal{F}_Z$ is slightly above the value at the injection point $2\mathcal{F}_s$, which is generally expected to be true for the maximum, due to noise fluctuations.

B. Search outliers of unknown origin

The S5R5 search additionally yielded eight search outliers of unknown origin above $\overline{2\mathcal{F}} \geq 6.5$. The results of the follow-up are summarized in Table III. None of these search outliers were found to be consistent with the signal hypothesis in the sense of Eq. (11): either they failed the \mathcal{F} -statistic consistency veto of Eq. (8) after refinement, or they were found to be consistent with Gaussian noise [in the sense of Eq. (9)] after the zoom stage.

These search outliers, with frequencies at approximately 434, 677, and ~ 984 Hz, are shown in Fig. 2 against the distribution of $\overline{2\mathcal{F}_Z}$ values obtained in Gaussian noise.

TABLE II. Most significant outlier after follow-up for each of the three fake pulsars. The injected signal parameters are f_s , α_s , δ_s , and \dot{f}_s . The value of the \mathcal{F} -statistic at the injection point is denoted as $\widetilde{2\mathcal{F}}_s$. The localization error of the final outlier in frequency and spin-down is Δf_Z and $\Delta \dot{f}_Z$, respectively, and $\Delta\gamma_Z = \arccos(\mathbf{n}_c \mathbf{n}_s)$ denotes the angular separation.

Fake Pulsar	Pulsar 2	Pulsar 3	Pulsar 5
f_s [Hz]	575.16355763140	108.857159397497	52.8083243593
α_s [rad]	3.75692884	3.11318871	5.28183129
δ_s [rad]	0.06010895	-0.58357880	-1.46326903
\dot{f}_s [Hz/s]	-1.37×10^{-13}	-1.46×10^{-17}	-4.03×10^{-18}
$\widetilde{2\mathcal{F}}_{\text{HS}}$	28	339	6.3
$\widetilde{2\mathcal{F}}_{\text{R}}$	100	1137	12
$\widetilde{2\mathcal{F}}_{\text{R}}^{\text{H1}}$	51	641	8.2
$\widetilde{2\mathcal{F}}_{\text{R}}^{\text{L1}}$	54	510	8.0
μ_{R}^*	4.01×10^{-4}	5.18×10^{-4}	4.88×10^{-3}
f_Z [Hz]	575.16355763214	108.857159397523	52.8083243548
α_Z [rad]	3.75692887	3.11318900	5.28181148
δ_Z [rad]	0.06010925	-0.58357884	-1.46326569
\dot{f}_Z [Hz/s]	-1.37×10^{-13}	3.30×10^{-16}	1.85×10^{-15}
$\widetilde{2\mathcal{F}}_Z$	7399	87097	678
$\widetilde{2\mathcal{F}}_Z^{\text{H1}}$	3519	47572	350
$\widetilde{2\mathcal{F}}_Z^{\text{L1}}$	3896	39557	332
$\widetilde{2\mathcal{F}}_s$	7377	86968	677
μ_Z^*	2.4×10^{-3}	1.4×10^{-3}	6.7×10^{-3}
Δf_Z [Hz]	7.44×10^{-10}	2.58×10^{-11}	-4.55×10^{-9}
$\Delta \dot{f}_Z$ [Hz/s]	-4.33×10^{-16}	3.44×10^{-16}	1.85×10^{-15}
$\Delta\gamma_Z$ [rad]	2.97×10^{-7}	2.48×10^{-7}	3.96×10^{-6}

TABLE III. Summary of the follow-up results for the three loudest search outliers associated with hardware injections and the eight most significant remaining outliers from the S5R5 search. The last column gives the classification of the final outlier after zoom, provided it passed the \mathcal{F} -statistic consistency veto.

f [Hz]	$\widetilde{2\mathcal{F}}_{\text{R}}$	$\widetilde{2\mathcal{F}}_{\text{R}}^{\text{H1}}$	$\widetilde{2\mathcal{F}}_{\text{R}}^{\text{L1}}$	\mathcal{F} -veto	$\widetilde{2\mathcal{F}}_Z$	$\widetilde{2\mathcal{F}}_Z^{\text{H1}}$	$\widetilde{2\mathcal{F}}_Z^{\text{L1}}$	Outcome
52	12	8.2	8.0	pass	678	350	332	S
96	9.1	4.4	13	fail
108	1137	641	510	pass	87097	47572	39557	S
144	11	4.5	14	fail
434	5.5	5.4	4.5	pass	47	30	22	G
575	100	51	54	pass	7399	3519	3896	S
677	6.4	5.4	5.2	pass	54	44	14	G
932	7.6	8.0	4.2	fail
984	6.5	4.8	5.5	pass	55	36	20	G
1030	7.4	8.3	4.5	fail
1142	8.5	10	4.2	fail

VI. DISCUSSION

In this paper we describe the extension of the two-stage follow-up method of [8] that was developed in order to follow up on search outliers from the Hough S5 Einstein@Home all-sky search [9]. The extension consists

of an additional Hough search as a prerefinement step, and an \mathcal{F} -statistic consistency veto after refinement to reduce the false-alarm rate on real detector data. Prerefinement was found to be necessary to improve the localization accuracy of the original search outliers.

With a Monte Carlo study, we quantify the detection probability as a function of the initial candidate strength, as shown in Fig. 3. In particular, we find that the pipeline achieves a detection probability of 90% for candidates without second-order spin-down at a strength of $\overline{2\mathcal{F}}_c \gtrsim 6$. On the other hand, for signals with maximal second-order spin-down (as considered by the original Hough Einstein@Home search [9]), the detection efficiency is reduced: for example, at $\overline{2\mathcal{F}}_c \approx 6.5$ the probability of signal recovery drops to $\approx 60\%$, while the pipeline is still able to separate $\approx 90\%$ of injected signals from Gaussian noise.

We illustrate the performance of this pipeline on real data by detailing the follow-up of Hough Einstein@Home search outliers, which was first presented in [9]. The pipeline successfully detects the three hardware injections present in the search outliers set and recovers their parameters with high accuracy (see Table II). The follow-up of the eight most significant search outliers of

unknown origin finds them to be consistent with either Gaussian noise or with line disturbances in the data.

ACKNOWLEDGMENTS

We are thankful for numerous discussions and comments from colleagues, in particular Badri Krishnan, Alicia Sintes, Bruce Allen, David Keitel, Karl Wette, and Stephen Fairhurst. We are grateful to Peter Shawhan, Teviet Creighton, and Andrzej Krolak for comments on this work in the process of review of [9]. M. S. gratefully acknowledges the support of Bruce Allen and the IMPRS on Gravitational Wave Astronomy of the Max Planck Society. P.L. and M.A.P. acknowledge support by the ‘‘Sonderforschungsbereich’’ Collaborative Research Centre (SFB/TR7). This work has been partly supported by NSF Grant No. PHY-1104902. This paper has been assigned AEI Preprint No. AEI-2014-009 and LIGO Document No. LIGO-P1400057-v3.

-
- [1] P. R. Brady and T. Creighton, *Phys. Rev. D* **61**, 082001 (2000).
 - [2] C. Cutler, I. Gholami, and B. Krishnan, *Phys. Rev. D* **72**, 042004 (2005).
 - [3] H. J. Pletsch and B. Allen, *Phys. Rev. Lett.* **103**, 181102 (2009).
 - [4] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, *Phys. Rev. D* **57**, 2101 (1998).
 - [5] P. Jaranowski, A. Krolak, and B. F. Schutz, *Phys. Rev. D* **58**, 063001 (1998).
 - [6] C. Cutler and B. F. Schutz, *Phys. Rev. D* **72**, 063006 (2005).
 - [7] C. Audet and J. E. Dennis, Jr., *SIAM J. Optim.* **17**, 188 (2006).
 - [8] M. Shaltev and R. Prix, *Phys. Rev. D* **87**, 084057 (2013).
 - [9] J. Aasi *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. D* **87**, 042001 (2013).
 - [10] B. Krishnan, A. Sintes, M. Papa, B. Schutz, S. Frasca, and C. Palomba, *Phys. Rev. D* **70**, 082001 (2004).
 - [11] R. Prix, *Phys. Rev. D* **75**, 023004 (2007).
 - [12] R. Prix and M. Shaltev, *Phys. Rev. D* **85**, 084010 (2012).
 - [13] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar, *Phys. Rev. D* **53**, 3033 (1996).
 - [14] B. J. Owen, *Phys. Rev. D* **53**, 6749 (1996).
 - [15] H. J. Pletsch, *Phys. Rev. D* **82**, 042002 (2010).
 - [16] D. Keitel, R. Prix, M. A. Papa, P. Leaci, and M. Siddiqi, *Phys. Rev. D* **89**, 064023 (2014).
 - [17] LALSUITE, <https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html> (2011).
 - [18] S. Le Digabel, *ACM Trans. Math. Softw.* **37**, 1 (2011).
 - [19] R. Prix, Technical Report No. LIGO-T0900149-v2, 2010.
 - [20] J. Aasi *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. D* **88**, 102002 (2013).
 - [21] B. Behnke, ‘‘A directed search for continuous gravitational waves from unknown isolated neutron stars at the galactic center,’’ <http://nbn-resolving.de/urn:nbn:de:gbv:089-7521728407> (2013).