Toward a computationally-efficient follow-up pipeline for blind continuous gravitational-wave searches

Lorenzo Mirasola [1, 2, *] and Rodrigo Tenorio [3, 4, 5, †]

[1] *Physics Department, Università degli Studi di Cagliari, Cagliari 09042, Italy*
[2] *INFN sezione di Cagliari, Cagliari 09042, Italy*
[3] *Dipartimento di Fisica "G. Occhialini", Università degli Studi di Milano-Bicocca, Piazza della Scienza 3, 20126 Milano, Italy*
[4] *INFN, Sezione di Milano-Bicocca, Piazza della Scienza 3, 20126 Milano, Italy*
[5] *Departament de Física, Universitat de les Illes Balears, IAC3 – IEEC, Carretera Valldemossa km 7.5, E-07122 Palma, Spain*

The sensitivity of continuous gravitational-wave (CW) searches for unknown neutron stars (NSs) is limited by their parameter space breadth. To fit within reasonable computing budgets, hierarchical schemes are used to identify interesting candidates using affordable methods. The resulting sensitivity depends on the number of candidates selected to follow-up. In this work, we present a novel framework to evaluate the effectiveness of stochastic CW follow-ups. Our results allow for a significant reduction of the computing cost of `pyfstat`, a well-established follow-up method. We also simplify the setup of multistage follow-ups by removing the need for parameter-space metrics. The study was conducted on Gaussian and real O3 Advanced LIGO data covering both isolated and binary sources. These results will have a positive impact on the sensitivity of all-sky searches in the forthcoming observing runs of the LIGO-Virgo-KAGRA collaboration.

I. INTRODUCTION

Continuous gravitational waves (CWs) are long-lasting gravitational-wave signals expected to be produced by rapidly spinning nonaxisymmetric neutron stars (NSs) [1]. Although currently undetected, CWs will extend our knowledge on galactic NSs, which so far have only been observed electromagnetically [2, 3].

All-sky searches attempt to detect CW signals from unknown sources. The resulting parameter-space breadth makes fully coherent matched filtering of year-long datasets computationally unfeasible [4]; semicoherent methods, which allow for coarser parameter-space discretizations, are used instead. Interesting candidates are then followed up using more sensitive methods [5]. These hierarchical schemes make all-sky searches more robust to stochastic NS physics, such as glitches [6] or accretion-induced spin wandering [7].

A favorable strategy to increase the sensitivity of a search without imposing stricter constraints on the signal model is increasing the number of candidates to follow up: generally, the more candidates are selected, the weaker a signal can be before being dismissed. As a result, all-sky search sensitivity is a monotonic function of the available follow-up computing budget. It is therefore fundamental to develop computationally efficient follow-up strategies to achieve a first CW detection from an unknown source.

In this work, we develop a novel framework to characterize the effectiveness of a generic CW follow-up as a function of signal mismatch and computing cost. The results of this study allow us to justify a five to ten-fold computing-cost reduction of a generic follow-up stage using `ptemcee` [8, 9] as implemented in `pyfstat` [10, 11], a state-of-the-art generic CW follow-up package based on Markov Chain Monte Carlo (MCMC) methods. Additionally, we simplify the setup of multistage follow-up schemes by removing the need for a full parameter-space metric [6].

This framework is applied to simulated all-sky searches for unknown NSs, both isolated and in binary systems, using simulated Gaussian noise and real data from the O3 LIGO-Virgo-KAGRA observing run [12–17]. We demonstrate the impact of our developments on real data by estimating the sensitivity improvement on an all-sky search in Advanced LIGO O3 data [18] due to the selection of a higher number of candidates.

The low-cost follow-up strategy discussed here is especially applicable to short-coherence CW searches, such as those described in [19–21], which operate by identifying narrow parameter-space regions of interest on which to deploy a more sensitive method. In this sense, this work complements that of [22], where a different follow-up approach, targeting broader parameter-space regions, is discussed.

The paper is structured as follows: in Sec. II, we revisit the basics of CW MCMC follow-ups and discuss a new approach to evaluate their sensitivity. In Sec. III, we formulate a multiobjective optimization problem to construct effective MCMC configurations. The effectiveness of different follow-up configurations is evaluated in Sec. IV by means of an injection campaign on Gaussian and real O3 data. The sensitivity impact of these improvements on a real search is discussed in Sec. V. We conclude in Sec. VI.

II. FOLLOW-UP OF CONTINUOUS-WAVE CANDIDATES

Blind searches identify interesting CW candidates using a computationally affordable method to then follow them up using a more sensitive approach. The purpose of

* lorenzo.mirasola@ca.infn.it
† rodrigo.tenorio@unimib.it

this hierarchical approach is to gradually reduce the false-alarm probability associated with a CW candidate while maintaining the parameter-space size under control [19, 23–25].

Several approaches have been proposed to follow-up CW candidates using long coherence times [20, 26–32]. These make use of lattice template banks, which are feasible for isolated sources but are impractical for sources in binary systems [33]. Lately, however, MCMC strategies have proven successful to follow-up candidates produced by a broad class of CW searches [10, 22, 34].

In this section, we revisit the fundamentals of stochastic multistage CW follow-up analyses of narrow parameter-space regions. The semicoherent $\mathcal{F}$-statistic and the basics of single-stage MCMC follow-ups are introduced in Secs. II A and II B. In Sec. II C we discuss, for the first time, the sensitivity of a follow-up in terms of its false-alarm and false-dismissal probabilities. In Sec. II D we propose a novel and simpler strategy to construct multistage follow-up schemes avoiding the use of a full parameter-space metric.

A. CW prolegomena

CW signals as emitted by rapidly spinning nonaxisymmetric NSs can be described using two sets of parameters: the *amplitude* parameters $\mathcal{A}$ and the *phase-evolution* parameters λ. This split is motivated by the detector's response to such signals

$$h(t;\lambda,\mathcal{A}) = \sum_{\mu=0}^{3} \mathcal{A}^{\mu}\, h_{\mu}(t;\lambda)\,, \tag{1}$$

where t refers to the time at the GW detector or Solar System Barycenter (SSB), depending on the specific pipeline implementation, and $\mathcal{A}^{\mu} = \mathcal{A}^{\mu}(\mathcal{A})$ are time-independent functions of the amplitude parameters that can be treated analytically [35–38]. In the case of a NS sustaining a quadrupolar deformation, the amplitude parameters include the initial phase ϕ_0, the polarization angle ψ, the cosine of the inclination angle with respect to the line of sight $\cos\iota$ and the nominal amplitude h_0, which depends on the specific CW emission mechanism [4].

For an isolated NS, the phase-evolution parameters λ_{i} include the GW frequency f_0, the linear spindown parameter f_1 at a fiducial reference time t_{ref} (further spindown terms can be included depending on the age of the source), and the sky position, which we parameterize using the equatorial angles $\hat{n} = (\alpha, \delta)$. The instantaneous frequency as measured by a detector is

$$f(t;\lambda_{\mathrm{i}}) = [f_0 + f_1(t - t_{\mathrm{ref}})] \cdot \left(1 + \frac{\vec{v}(t)}{c}\cdot\hat{n}\right)\,, \tag{2}$$

where $\vec{v}$ is the velocity of the detector with respect to the SSB.

For a NS in a circular binary system, phase-evolution parameters λ_{b} include the orbital period P or orbital frequency $\Omega = 2\pi/P$, the projected semimajor axis (in light seconds) a_{p}, and the time of passage through the ascending node t_{asc} [39]:

$$f(t;\lambda_{\mathrm{b}}) = f_0 \cdot \left(1 + \frac{\vec{v}(t)}{c}\cdot\hat{n} - a_{\mathrm{p}}\Omega\cos\left[\Omega(t - t_{\mathrm{asc}})\right]\right)\,. \tag{3}$$

These searches tend to omit spindown terms and eccentricity as their effect on a circular-orbit template is unresolvable for typical search stages [33, 40, 41].

Most all-sky CW searches operate by comparing the detector data x to a bank of waveform templates $\{\lambda_j, j = 1,\ldots,N_{\mathrm{T}}\}$ by means of a detection statistic. The long duration of these signals makes fully coherent matched filtering unaffordable; instead, these searches make use of semicoherent methods, which split the data into N_{seg} segments with a certain duration T_{coh} from which a detection statistic is computed (see [4] for a recent review).

In a Bayesian context [37, 42], semicoherent detection statistics can be related to a targeted marginal Bayes factor $B_{\mathrm{SG}}(x;\lambda)$ comparing the Gaussian noise hypothesis $\mathcal{H}_{\mathrm{G}}$, under which the data consist of Gaussian noise samples $x = n$, versus the signal hypothesis $\mathcal{H}_{\mathrm{S}}(\lambda,\mathcal{A})$, which includes a CW signal with a specific set of parameters $x = n + h(\lambda,\mathcal{A})$:

$$B_{\mathrm{SG}}(x;\lambda) = \int \mathrm{d}\mathcal{A}\, \frac{p(x|\mathcal{H}_{\mathrm{S}}(\lambda,\mathcal{A}))}{p(x|\mathcal{H}_{\mathrm{G}})} p(\mathcal{A})\,. \tag{4}$$

(Note that we marginalize the likelihood *ratio* rather than the signal-hypothesis likelihood itself as $\mathcal{H}_{\mathrm{G}}$ is independent of $\mathcal{A}$). By splitting the dataset into disjoint segments and independently marginalizing the amplitude parameters in each of them we obtain the semicoherent $\mathcal{F}$-statistic [35, 36, 42]

$$B_{\mathrm{SG}}(x;\lambda) \propto e^{\hat{\mathcal{F}}(\lambda;x)}\,, \tag{5}$$

$$\hat{\mathcal{F}}(\lambda;x) = \sum_{s=1}^{N_{\mathrm{seg}}} \mathcal{F}_s(\lambda;x)\,, \tag{6}$$

where $\mathcal{F}_s$ refers to the fully coherent $\mathcal{F}$-statistic computed on segment s. Intuitively, splitting the data into segments trades "height" for "breadth", in the sense that the $\mathcal{F}$-statistic becomes more permissive to parameter mismatches (i.e. the peak becomes "broader" in the parameter space) at the price of shifting the background upward [5, 10, 28, 43, 44]. This same mechanism explains why increasing the coherence time is a sensible method to increase the significance of a candidate.

B. Single MCMC follow-ups using `pyfstat`

The basic idea in [10] is that, for a given dataset x, the *posterior probability* on the phase-evolution parameters λ can be expressed as

$$p(\lambda|x,\mathcal{H}_{\mathrm{S}}(\lambda,\mathcal{A})) \propto B_{\mathrm{SG}}(x;\lambda)\, p(\lambda)\,, \tag{7}$$

where $p(\lambda)$ corresponds to the prior distribution, which in this context corresponds to the uncertainty on the candidate's parameters produced by a search. It follows then that an adaptive template bank around the true signal parameters can be efficiently constructed by sampling the posterior distribution using the semicoherent $\mathcal{F}$-statistic. In this work, we consider $p(\lambda)$ to be a narrow prior with a width comparable to the typical parameter-space resolution of a search [6, 34]; this choice has proven successful in the follow-up of candidates from searches using relatively short coherence times [40, 45–50]. We refer the reader to [22] for the practicalities of following-up broader parameter-space regions.

As implemented in `pyfstat` [10, 11], the sampling of Eq. (7) is performed using `ptemcee`, a parallel-tempering ensemble MCMC sampler [8, 9], and the $\mathcal{F}$-statistic uses the tools provided by LALSuite [51, 52]. `ptemcee` takes in three configuration parameters: the total number of steps $N_{\rm tot}$, the number of walkers in the ensemble $N_{\rm w}$, and the number of parallel temperatures $N_{\rm t}$. The total computing cost of a configuration $c = (N_{\rm tot}, N_{\rm w}, N_{\rm t})$, in units of $\mathcal{F}$-statistic evaluations, can be computed as [10]

$$\mathcal{C}(c) = N_{\rm tot} N_{\rm t} N_{\rm w} \,. \tag{8}$$

We postpone to Sec. III the selection of an appropriate configuration by considering the false-alarm and false-dismissal probabilities of a follow-up stage, which we derive next in Sec. II C.

C. Expected sensitivity of a follow-up stage

The sensitivity of a CW all-sky search is typically reported as an estimate of the required amplitude for a population of isotropically oriented all-sky distributed NSs so that a significant fraction of them (usually 90% or 95%) is deemed as "detected" [53]. These populations can be labeled using the sensitivity depth [54, 55]

$$\mathcal{D} = \frac{\sqrt{S_{\rm n}}}{h_0} \,, \tag{9}$$

where $S_{\rm n}$ represents the single-sided power spectral density of the noise (PSD).

A single-stage follow-up evaluates the semicoherent $\mathcal{F}$-statistic on a finite number of templates randomly sampled by an MCMC. Upon completion, the $\mathcal{F}$-statistic of the loudest template is compared against a threshold $2\mathcal{F}_{\rm thr}$ to decide whether to discard the candidate or carry it on to the next stage. If the follow-up is part of a broader set of stages, this veto may depend on results obtained by previous stages, e.g. [18, 32, 34, 49, 56]. To set up an effective follow-up, it is critical to understand the required detectability thresholds.

The sensitivity of an $\mathcal{F}$-statistic search was thoroughly discussed in [53, 55]. Here we reintroduce the basic principles and release a new Python code, `cows3` [57], which implements a version of their semianalytical method. We also extend the discussion to consider the false-alarm probability associated to an MCMC follow-up in Sec. II C 2.

1. *Sensitivity depth at fixed false-dismissal probability*

The standard use-case of `pyfstat`-like follow-ups [18, 33, 40, 41, 45–50] calibrates a threshold such that only a "small" fraction of the signal population is dismissed. This fraction is quantified by means of the false-dismissal probability $p_{\rm fd}$, and is a function of both the number of segments $N_{\rm seg}$ (or coherence time $T_{\rm coh}$) and sensitivity depth of the population. Understanding the relation between $p_{\rm fd}$, $\mathcal{D}$, and $2\mathcal{F}_{\rm thr}$ is thus required to set up an effective follow-up. Throughout this derivation, we assume $S_{\rm n}$ has been properly estimated [19, 58, 59] and keep it implicit in our notation.

The false-dismissal probability associated to a threshold $2\mathcal{F}_{\rm thr}$ under $\mathcal{H}_{\rm S}(\lambda, \mathcal{A})$ is given by

$$p_{\rm fd}(2\mathcal{F}_{\rm thr}, \mathcal{D}) = p(2\hat{\mathcal{F}} < 2\mathcal{F}_{\rm thr}|\mathcal{D}) = \int_{-\infty}^{2\mathcal{F}_{\rm thr}} {\rm d}2\hat{\mathcal{F}}\, p(2\hat{\mathcal{F}}|\mathcal{D}) \,, \tag{10}$$

where, to ease the notation we have omitted the number of coherent segments $N_{\rm seg}$ and $\mathcal{H}_{\rm S}(\lambda, \mathcal{A})$ will be replaced by $\mathcal{D}$ or $\mathcal{A}$ whenever convenient. Here, $p(2\hat{\mathcal{F}}|\mathcal{D})$ represents the $\mathcal{F}$-statistic distribution associated to a matched (cf. mismatch discussion in Sec. II C 3) CW signal drawn from a population at depth $\mathcal{D}$, which is inconvenient to evaluate and depends on the source distribution of amplitude parameters and sky positions through a marginalization integral

$$p(2\hat{\mathcal{F}}|\mathcal{D}) = \int {\rm d}\mathcal{A}\, {\rm d}\hat{n}\, p(2\hat{\mathcal{F}}|\mathcal{A}\,\hat{n})\, p(\mathcal{A}\,\hat{n}|\mathcal{D}) \,. \tag{11}$$

where $p(2\hat{\mathcal{F}}|\mathcal{A}\,\hat{n}\,\mathcal{D}) = p(2\hat{\mathcal{F}}|\mathcal{A}\,\hat{n})$, as $\mathcal{A}$ and $S_{\rm n}$ are enough to compute $\mathcal{D}$.

For a given observing run, the $\mathcal{F}$-statistic depends on the signal's sky position $\hat{n}$ and amplitude parameters $\mathcal{A}$ through a deterministic quantity $\rho^2(\mathcal{A}, \hat{n})$, which is usually referred to as the (squared) signal-to-noise-ratio (SNR, see [55] for an explicit expression for ρ^2). This allows for Eq. (11) to be expressed as a one-dimensional integral

$$p(2\hat{\mathcal{F}}|\mathcal{D}) = \int_0^\infty {\rm d}\rho^2 p(2\hat{\mathcal{F}}|\rho^2) p(\rho^2|\mathcal{D}) \,, \tag{12}$$

where

$$p(\rho^2|\mathcal{D}) = \int {\rm d}\mathcal{A}{\rm d}\hat{n}\, p(\rho^2|\mathcal{A}\,\hat{n}) p(\mathcal{A}\,\hat{n}|\mathcal{D}) \tag{13}$$

and $p(2\hat{\mathcal{F}}|\rho^2) = \chi^2_{4N_{\rm seg}}(2\hat{\mathcal{F}}; \rho^2)$ is a noncentral chi-squared distribution with $4N_{\rm seg}$ degrees of freedom and noncentrality parameter ρ^2 [35, 36]. Moreover, the sensitivity depth can be factored out of the SNR as

$$\rho^2(\mathcal{A}, \hat{n}) = \frac{1}{\mathcal{D}^2} \rho_0^2(\mathcal{A}_0, \hat{n}) \,, \tag{14}$$

where $\mathcal{A}_0 = \{\cos\iota, \psi\}$[1]. Equation (10) is thus an expectation over the distribution of ρ_0^2

$$p_{\rm fd}(2\mathcal{F}_{\rm thr}, \mathcal{D}) = \int_0^\infty {\rm d}\rho_0^2\, p(\rho_0^2) p(2\hat{\mathcal{F}} < 2\mathcal{F}_{\rm thr} | \rho^2 = \rho_0^2/\mathcal{D}^2)\,, \tag{15}$$

where

$$p(\rho_0^2) = \int {\rm d}\mathcal{A}_0 {\rm d}\hat{n}\, p(\rho_0^2|\mathcal{A}_0\,\hat{n})\, p(\mathcal{A}_0\,\hat{n}) \tag{16}$$

and $p(\rho_0^2|\mathcal{A}_0\,\hat{n}) = \delta(\rho_0^2 - \rho_0^2(\mathcal{A}_0, \hat{n}))$.

The key insight in [53, 55] is that, once the signal population has been specified through the prior $p(\mathcal{A}_0, \hat{n})$, Eq. (16) is *independent of the signal's amplitude* and only needs to be generated *once* per observing run and signal population. As a result, $p(\rho_0^2)$ can be numerically generated and stored as a one-dimensional histogram with which Eq. (15) can be easily computed for a given $2\mathcal{F}_{\rm thr}$ and $\mathcal{D}$ [2].

In Fig. 1 we compute the false-dismissal probability using Eq. (15) for three representative depth values, namely $\mathcal{D} = \{20, 30, 50\}/\sqrt{\rm Hz}$. These cover the typical range of sensitivities achieved by all-sky searches with $T_{\rm coh} \lesssim 0.5\,$days [18, 30, 40, 41, 61–63]. We simulated a 1-year observing run with two detectors consistent with the Advanced LIGO detectors [64], LIGO Hanford (H1) and LIGO Livingston (L1), starting at the beginning O3. Threshold values are computed using $N_{\rm seg} = 730$ (i.e. $T_{\rm coh} = 0.5\,$d), which is a standard choice for a first-stage follow-up [18, 40, 45, 46, 48, 65, 66].

Generally, single-stage follow-ups estimate $p_{\rm fd}$ as the fraction of dismissed signals in an simulated injection campaign; as a result, for a given search, $p_{\rm fd}$ is lower-bounded by the number of simulated signals. By means of Eq. (15), on the other hand, we can compute arbitrarily low false-dismissal probabilities (insofar the signal model in use matches the expected signal). Typical follow-ups operate at $p_{\rm fd} \in [10^{-5}, 10^{-3}]$ [18, 33, 40, 41]. In this work, we chose a conservative value of $p_{\rm fd} = 10^{-5}$ to present our results.

2. *A note on $p_{\rm fa}$*

Given a threshold $2\mathcal{F}_{\rm thr}$, the corresponding false-alarm probability $p_{\rm fa}$ can be evaluated as the survival function of the detection statistic at hand under the noise hypothesis

$$p_{\rm fa}(2\mathcal{F}_{\rm thr}) = \int_{2\mathcal{F}_{\rm thr}}^\infty {\rm d}2\hat{\mathcal{F}}\, p(2\hat{\mathcal{F}}|\mathcal{H}_{\rm G})\,. \tag{17}$$

[1] $\rho_0^2(\mathcal{A}_0, \hat{n})$, which can be interpreted as a "unit-depth SNR", will play a similar role to the "response function" $R^2(\theta)$ in [53, 55].

[2] Incidentally, given a desired false-dismissal probability $p_{\rm fd}^*$, we can seamlessly estimate $\mathcal{D}(2\mathcal{F}_{\rm thr}, p_{\rm fd}^*)$ by using a root-finding algorithm on Eq. (15), as implemented in [60]. We provide the first open-source Python implementation of this algorithm in [57].

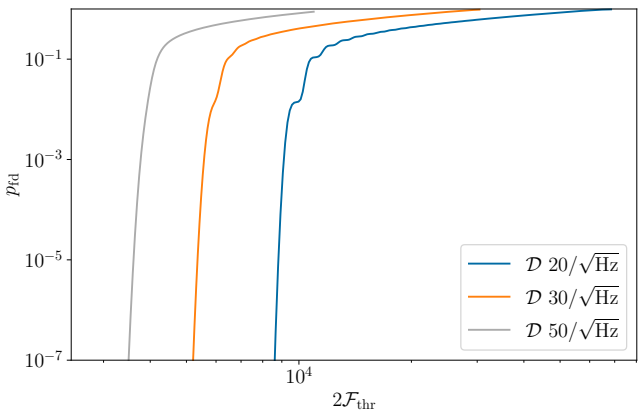

FIG. 1. False-dismissal probabilities [Eq. (15)] for a simulated 1-year observing run with two detectors compatible with H1 and L1 using $N_{\rm seg} = 730$ ($T_{\rm coh} = 0.5\,$d) for three representative sensitivity depths.

This quantity could be interpreted as the fraction of nonastrophysical CW candidates that would not be discarded by a follow-up stage assuming Gaussian noise. Note, however, that real searches are affected by non-Gaussian noise and disturbances implying that a larger fraction of non-CW-related outliers might pass this first selection. Moreover, we here neglect the fact that selected candidates in a follow-up correspond to the *loudest* candidates of a previous search step.

The distribution of the $\mathcal{F}$-statistic under the Gaussian noise hypothesis is well-known and corresponds to a central chi-squared distribution with $4N_{\rm seg}$ degrees of freedom [35, 36]. In a follow-up, similarly to some searches (e.g. [67–70]), the statistic upon which a threshold is imposed is actually the $\mathcal{F}$-statistic of the *loudest* template, $2\hat{\mathcal{F}}_{\rm max}$. The distribution of such statistic was discussed in [34, 69], and corresponds to a Gumbel distribution

$$p(2\hat{\mathcal{F}}_{\rm max}|\mathcal{H}_{\rm G}) = \frac{1}{\sigma_{\rm G}} e^{-\left(\frac{2\hat{\mathcal{F}}_{\rm max}-\mu_{\rm G}}{\sigma_{\rm G}}\right) - e^{-\left(\frac{2\hat{\mathcal{F}}_{\rm max}-\mu_{\rm G}}{\sigma_{\rm G}}\right)}}\,, \tag{18}$$

where the location $\mu_{\rm G}$ and scale $\sigma_{\rm G}$ parameters can be numerically characterized. As discussed in [69], Eq. (18) is generally applicable in the presence of non-Gaussianities, such as spectral lines [22, 71].

It was first proposed in [34] to use off-sourcing [72] to re-evaluate the template bank produced by an MCMC follow-up at different sky points and get the loudest point to numerically generate $p(2\hat{\mathcal{F}}_{\rm max}|\mathcal{H}_{\rm G})$. During the development of this work, we noted that such an approach *underestimates* the expected background of an MCMC. This is related to the fact that MCMC sampling tends to favor high-likelihood regions, as opposed to a random draw in which different templates are drawn independently. The loudest value of an MCMC, thus, *tends to be higher* than the loudest of a random template

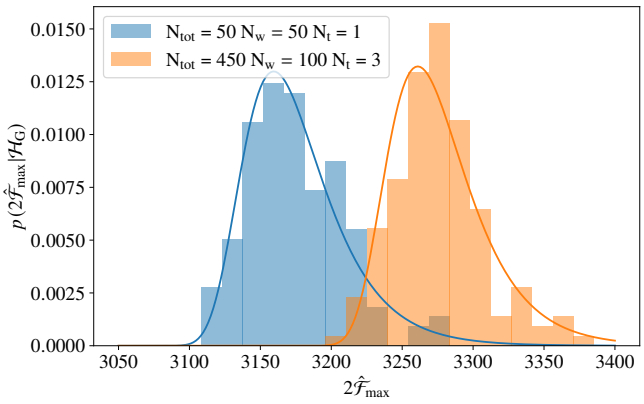

FIG. 2. Noise distributions for two particular follow-up configurations, as specified in the legend. Each histogram is generated using 150 Gaussian noise realizations of 1-year observing runs with 2 detectors using $N_{\rm seg} = 730$. The solid lines represent the best fit for a Gumbel distribution using the method of `scipy.stats.gumbel_r` [73]. Results are obtained using the priors for a search of isolated emitters in Table II with width calculated for $T_{\rm coh} = T_{\rm SFT}$. LM: Lorenzo remove green histo

bank[3].

In Fig. 2, we numerically generate the noise distribution of two particular MCMC configurations with the priors listed in Table II for a search of isolated emitters. We generated 150 Gaussian noise realizations with the same data setup as in Fig. 1. This number of realizations allows us to estimate false-alarm probabilities down to $p_{\rm fa} \gtrsim 10^{-2}$, which is enough for a first-stage follow-up (see e.g. [18]). We observe that the general behavior of the distribution is well captured by a Gumbel distribution, as discussed in [69]. We also note that the specific parameters of the Gumbel distribution are a function of the chosen configuration. To draw fair conclusions, we estimate the Gumbel parameters for a relevant subset of the MCMC considered in this work. The results are collected in Appendix A.

3. *Mismatch corresponds to false-alarm probability*

Throughout Sec. II C 1 we implicitly assumed the phase-evolution parameters of a candidate were perfectly matched by the follow-up's template bank, so that the recovered SNR corresponded to the actual SNR of the signal. In practice, however, the signal's parameters $\lambda_{\rm s}$ will be off by $\Delta\lambda$ from the closest template in the bank. This deviation is quantified by *mismatch*, which is defined as the fractional SNR loss due to $\Delta\lambda$ [74]

$$\mu(\Delta\lambda;\lambda_{\rm s}) = \frac{2\hat{\mathcal{F}}(\lambda_{\rm s}) - 2\hat{\mathcal{F}}(\lambda_{\rm s}+\Delta\lambda)}{2\hat{\mathcal{F}}(\lambda_{\rm s}) - 4N_{\rm seg}}\,, \tag{19}$$

where $4N_{\rm seg}$ corresponds to the expected $2\hat{\mathcal{F}}$ value in noise.

To account for a mismatch μ, the threshold $2\mathcal{F}_{\rm thr}$ at a fixed $p_{\rm fd}$ must be *lowered* to

$$2\mathcal{F}_{\rm thr}(\mu) = 2\mathcal{F}_{\rm thr}\cdot(1-\mu) + 4N_{\rm seg}\mu\,, \tag{20}$$

which corresponds to a reduction from the optimal SNR ρ^2 to $\rho^2(1-\mu)$. This implies that, at a given $p_{\rm fd}$, the false-alarm probability of the follow-up is now a function of mismatch

$$p_{\rm fa}(\mu) = \int_{2\mathcal{F}_{\rm thr}(\mu)}^{+\infty} {\rm d}2\hat{\mathcal{F}}\, p(2\hat{\mathcal{F}}_{\rm max}|\mathcal{H}_{\rm G})\,. \tag{21}$$

We show the resulting false alarm probability using the distributions discussed in Sec. II C 2 and Appendix A in Fig. 3. Depending on sensitivity depth, acceptable mismatches (i.e. $p_{\rm fa}(\mu) \approx 10^{-2}$) are as high as $\mu \approx 0.8$.

The result in Eq. (21) culminates in the proposal of mismatch μ as a metric for the effectiveness of a follow-up: *any* signal can be "detected" by a follow-up if enough false-alarm probability is allowed (i.e. a low-enough threshold is used). The goal of an effective follow-up is to claim said "detection" in such a way that *a significant fraction of the noise candidates* is discarded in a single stage. Failing to operate at a low false-alarm probability defeats the purpose of a fast follow-up, as computing cost is then simply pushed forward to further stages.

Note that the $p_{\rm fa}$ here computed is a simple construct to compare the effectiveness of different configurations on equal grounds. This is particularly important in the case of a follow-up, as the selected candidates are usually not a random selection but rather the *top-scoring* templates from a previous search stage, thus implying our $p_{\rm fa}$ could fall short on a real search case. However, the typical difference in segment lengths between a search such as the ones here considered ($N_{\rm seg} \sim 10^4$) and a follow-up ($N_{\rm seg} \sim 10^3$) is likely to diminish the effects of such choice. Since follow-ups tend to operate by selecting candidates based on computing cost and not false-alarm probability, we leave such a discussion for future work and limit the usage of $p_{\rm fa}$ to the comparison of different MCMC configurations.

D. A new coherence-time ladder

As a final step, we address the setup of a coherence-time ladder. As described in [6, 34], this is a multistage follow-up setup in which coherence time (number of segments) is gradually increased (decreased) in order to impose stronger constraints and zoom in on CW candidates. This can be seen as a rudimentary form of sequential Monte Carlo [75].

[3] This has no effect on past results, as every application calibrated its own follow-up. On the other hand, this result explains the unusually high Bayes factors quoted in [34], where the noise distribution was severely underestimated.

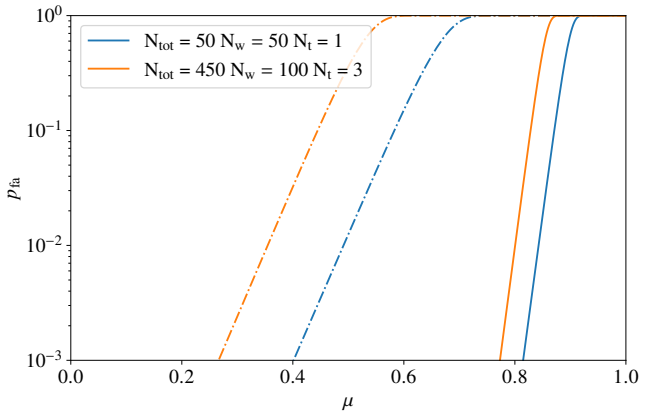

FIG. 3. False-alarm probability as a function of mismatch for the two MCMC configurations shown in Fig. 2. Signal populations are compatible with a population of isolated neutron stars at $\mathcal{D} = 30/\sqrt{\mathrm{Hz}}$ (solid lines) and $\mathcal{D} = 50/\sqrt{\mathrm{Hz}}$ (dash-dotted lines).

The standard construction of the ladder [6] uses the parameter-space metric, which is only known analytically for a limited number of cases [23, 39, 76–83]; for more general cases, a numerical approach is required, as discussed in [22]. Here, we propose a simpler approach using the scaling of individual parameter-space dimensions with T_{coh}.

The ladder construction goes as follows: suppose we conduct a follow-up on a parameter space region $\Delta\lambda^{(j)}$ using a coherence time $T^{(j)}_{\mathrm{coh}}$. The number of "independent templates" in said region can be estimated as

$$\mathcal{N}(\Delta\lambda^{(j)}, T^{(j)}_{\mathrm{coh}}) = \int_{\Delta\lambda^{(j)}} \mathrm{d}\lambda \sqrt{g\left(T^{(j)}_{\mathrm{coh}}\right)}\,, \tag{22}$$

where g is the determinant of the parameter-space metric. It was argued in [6] that `pyfstat` is effective for parameter-space regions containing $\mathcal{N}_*$ templates, typically around 10^3 to 10^4. This was reinterpreted in [34] as the fraction of volume down to which a follow-up can successfully compress a posterior distribution. Broader regions were later explored in [22].

After a successful follow-up stage, the resulting region $\Delta\lambda^{(j+1)}$ will contain a significantly smaller number of templates than the original one, $\mathcal{N}(\Delta\lambda^{(j+1)}, T^{(j)}_{\mathrm{coh}}) \sim 1$. Also, in this work we consider follow-ups as local analyses, which means the number of templates can be approximated as

$$\mathcal{N}(\Delta\lambda^{(j)}, T^{(j)}_{\mathrm{coh}}) \approx \mathrm{Vol}(\Delta\lambda^{(j)}) \sqrt{g\left(T^{(j)}_{\mathrm{coh}}\right)}\,. \tag{23}$$

The new coherence time $T^{(j+1)}_{\mathrm{coh}}$, thus, should be chosen such that

$$\begin{aligned}\mathcal{N}_* &= \mathcal{N}(\Delta\lambda^{(j+1)}, T^{(j+1)}_{\mathrm{coh}}) \\ &\approx \frac{\mathcal{N}\left(\Delta\lambda^{(j+1)}, T^{(j+1)}_{\mathrm{coh}}\right)}{\mathcal{N}\left(\Delta\lambda^{(j+1)}, T^{(j)}_{\mathrm{coh}}\right)} \approx \frac{\sqrt{g\left(T^{(j+1)}_{\mathrm{coh}}\right)}}{\sqrt{g\left(T^{(j)}_{\mathrm{coh}}\right)}}\,.\end{aligned} \tag{24}$$

where to go to the second line we used $\mathcal{N}(\Delta\lambda^{(j+1)}, T^{(j)}_{\mathrm{coh}}) \sim 1$ to introduce a term containing $\mathrm{Vol}(\Delta\lambda^{(j)})$ to cancel it from the numerator, as described in [6]. An alternative interpretation is given in [34].

Equation (24) requires to evaluate the determinant of the parameter-space metric. Since the follow-up is a local analysis, any dependency on the parameter-space position is likely to be negligible and coherence time T_{coh} will be the only quantity that changes across stages. We also note that, for semicoherent searches, the resolution of relevant phase-evolution parameters tends to scale with T^{-1}_{coh} [19, 39, 78]. These two facts motivate the following ansatz to estimate the number of templates in a follow-up stage

$$\mathcal{N}_* = \left(\frac{T^{(j+1)}_{\mathrm{coh}}}{T^{(j)}_{\mathrm{coh}}}\right)^D\,. \tag{25}$$

Here, D is the number of resolved parameters considered by the follow-up. Therefore, the coherence-time ladder can be constructed as easily as

$$T^{(j)}_{\mathrm{coh}} = \mathcal{N}_*^{j/D}\, T^{(0)}_{\mathrm{coh}}\,. \tag{26}$$

Choosing an appropriate value of D is not trivial in general, as different parameter-space dimensions will have different dependencies on T_{coh} depending on typical variation scale of the signal's phase (see e.g. the short-segment and long-segment metrics in [39]). In this work, we choose D assuming that all the resolvable parameters scale with T^{-1}_{coh}. This returns a higher value than required, and thus allows us to operate in the safe side. Since we are dealing with isolated signals with one spindown term and circular binaries, we choose $D = 4$ for the isolated case and $D = 6$ for the binary case. Other signal models, such as higher number of spindowns or eccentricity in the binary orbit, may require extra stages, depending on the chosen configuration, to ensure proper convergence.

III. HOW TO EVALUATE THE EFFECTIVENESS OF CW FOLLOW-UPS

We are interested in follow-up configurations with 1) a low computing cost and 2) a low false-alarm probability at a fixed false-dismissal probability and sensitivity depth. This is equivalent to a multiobjective optimization problem, which is in general solved by a *family of equivalent solutions*, each of them representing a different trade-off between 1) and 2) [84]. As per Sec. II,

rather than minimizing $p_{\rm fa}$ an effective follow-up can be defined so that $p_{\rm fa}$ is below a *maximum* allowed value; using Eq. (21), this corresponds to a maximum allowed mismatch $\mu_{\rm M}$. As a result, we define the optimal follow-up configuration c_* so that

$$c_* = \arg\min_c \mathcal{C}(c)\,, \mu(c) < \mu_{\rm M}\,, \tag{27}$$

where $\mathcal{C}(c)$ is computed using Eq. (8) and $\mu(c)$ is the mismatch associated to a follow-up configuration c.

The considered list of follow-up configurations is shown in Table I. We detail the computation of $\mu_{\rm M}$ in Sec. III A and $\mu(c)$ in Sec. III B. This process makes use of a campaign of injected software-simulated signals described in Sec. III C. Results are presented in Sec. IV.

A. Maximum allowed mismatch

As discussed in Sec. II C, for a given population depth $\mathcal{D}$ and false-dismissal probability $p_{\rm fd}$, and for the case $\mu = 0$, we can construct a detectability threshold $2\mathcal{F}_{\rm thr}$ by numerically inverting Eq. (10). To account for a mismatch μ, this threshold should be modified using Eq. (20) into $2\mathcal{F}_{\rm thr}(\mu)$. Finally, Eq. (21) links the detectability threshold $2\mathcal{F}_{\rm thr}(\mu)$ to the resulting false-alarm probability $p_{\rm fa}(\mu, c)$.[4] This establishes a monotonic relation between $p_{\rm fa}$ and μ (for a given $p_{\rm fd}$, $\mathcal{D}$, and c) which can be numerically inverted to compute $\mu_{\rm M} = \mu(p_{\rm fa}, c)$.

We consider the case $p_{\rm fd} = 10^{-5}$ and a range of false-alarm probabilities $p_{\rm fa} \in \left[10^{-2}, 10^{-1}\right]$. This produces a range of $\mu_{\rm M}$ values for each configuration c. To evaluate our results, we consider the broadest range of $\mu_{\rm M}$ across all configurations at a fixed cost, which we show as shaded areas in Figs. 4 and 5. Acceptable follow-up configurations are lower than the shaded area.

B. Mismatch of a configuration

To associate a mismatch to a configuration c, we use a software-simulated injection campaign, described in Sec. III C. For each injection, we run an MCMC follow-up using the configuration c and retrieve the loudest $\mathcal{F}$-statistic. This value is compared to the $\mathcal{F}$-statistic at the injection point to compute the resulting mismatch [Eq. (19)]. The result is a sample of mismatches, one for each injection, which together describe the mismatch distribution associated to the chosen MCMC configuration c. We then take the $90^{\rm th}$ quantile of this sample as a proxy for the mismatch of the configuration $\mu(c) = \mu^{90\%}(c)$. This last choice is motivated by the available computing resources; a subset of the configurations was re-evaluated on a larger injection campaign using the 95th and 99th percentile, with consistent results. All the injections used in this procedure return an $\mathcal{F}$-statistic above the detectability thresholds associated to the probed values of $p_{\rm fd}$ and $p_{\rm fa}$ in the first follow-up stage.

$N_{\rm w}$	50, 100
$N_{\rm t}$	1, 2, 3
$N_{\rm tot}$	50, 100, . . . , 450, 500

TABLE I. Parameter space of the MCMC configurations considered in this work. Each of the 54 configurations c is an element of the cartesian product of these three parameters. $N_{\rm tot}$ is varied in steps of 50.

We also compute the track distance [85]

$$d(\lambda, \lambda_{\rm s}) = T_{\rm SFT}\langle |f(t;\lambda) - f(t;\lambda_{\rm s})| \rangle\,, \tag{28}$$

which evaluates the distance between two templates as the average area enclosed between their corresponding frequency tracks on the spectrogram. The notation $\langle\cdot\rangle$ denotes averaging along a subset of the timestamps of the observing run, as discussed in [85]. This will allow us to assess whether the loudest candidate from an MCMC follow-up corresponds to the injected signal or a noise fluctuation.

C. Injection campaign

We simulate two injection campaigns of isotropically oriented all-sky CW sources in Gaussian and real noise in a consistent manner with the two all-sky CW searches produced by the LVK in O3a and O3 data [18, 40]. The resulting MCMC follow-up performances will thus be informative to future all-sky searches. In both cases, we consider two detectors in a location consistent with the Advanced LIGO detectors.

We assume a PSD of $\sqrt{S_{\rm n}} = 10^{-23}/\sqrt{\rm Hz}$ for gapless Gaussian noise. O3 data is generated using the science segments and gating times available in [86, 87]. The injection parameter spaces are summarized in Table II. Frequency and binary orbital parameters are drawn uniformly in the specified ranges; sky positions are drawn uniformly across the celestial sphere; spindown values are drawn from a mixture of two log-uniform distributions with relative weights of 0.9 for the negative range and 0.1 for the positive range; the cosines of the inclination angles, the polarization angles, and the initial phase are drawn from a uniform distribution along $[-1, 1]$, $[-\pi/4, \pi/4]$ and $[0, 2\pi]$, respectively.

For sources in binary systems, we use $T_{\rm obs} = 0.5$ year and generate 50%-overlapping Short Fourier Transforms (SFTs) [88] with a duration of $T_{\rm SFT} = 1024\,$s to be consistent with [40]. For isolated sources, we use $T_{\rm obs} = 1$ year and nonoverlapping $T_{\rm SFT} = 1800\,$s SFTs. Signals are injected at two representative sensitivity depths, namely 30 and $50/\sqrt{\rm Hz}$ for isolated and 20 and $25/\sqrt{\rm Hz}$ for sources in binary systems. This discrepancy

[4] Note that this last step depends on c through $p(2\hat{\mathcal{F}}_{\rm max}|\mathcal{H}_{\rm G})$.

between isolated and binary sources arises from the different parameter-space dimensionality.

For completion, the cost of a single $\mathcal{F}$-statistic evaluation, which depends on the number of SFTs [89], is about 0.4 s to 0.7 s in a 13th Gen. Intel(R) Core(TM) i7-1355U.

	Interval	Prior width
f_0	[50, 500] Hz	$\frac{1}{2T_{\rm coh}}$
f_1	$[-10^{-8}, 10^{-9}]$ Hz/s	$\frac{1}{2T_{\rm coh}T_{\rm obs}}$
α	$[-\pi, \pi]$ rad	$\frac{10^4}{2fT_{\rm coh}}$
δ	$[-\pi/2, \pi/2]$ rad	$\frac{10^4}{2fT_{\rm coh}}$
P	[7, 15] d	$\frac{\pi}{f\,a_p T_{\rm coh}T_{\rm obs}\Omega^3}$
a_p	[5, 15] lt-s	$\frac{1}{2fT_{\rm coh}\Omega}$
$t_{\rm asc}$	$[t_{\rm mid}-\frac{P}{2}, t_{\rm mid}+\frac{P}{2}]$	$\frac{5}{2f\,a_p\,T_{\rm coh}\,\Omega^2}$

TABLE II. Parameter-space covered by the injection campaigns. The prior's width corresponds to the standard deviation of a Gaussian prior, except for $t_{\rm asc}$ which uses a uniform prior. $t_{\rm mid}$ refers to the timestamp at the middle of the observing run. For the first follow-up stage, we calculate the widths using $T_{\rm coh} = T_{\rm SFT}$.

IV. OPTIMIZING THE CHOICE OF FOLLOW-UP SETUP

We present the results of evaluating the 54 different MCMC configurations listed in Table I on the injection campaigns described in Sec. III C. Additionally, we run an extra MCMC configuration corresponding to the standard choice in previous MCMC follow-ups [18, 40, 45, 46, 48, 65, 66] given by $N_{\rm w} = 100, N_{\rm t} = 3, N_{\rm tot} = 600$.

We use uncorrelated Gaussian priors with a width given by the typical parameter resolution (see Table II). This is consistent with a maximum-entropy distribution given the average location and uncertainty of a parameter [90]. To mimic the behavior of an actual search, we randomly shift the priors' centre away from the injection point. To do so, we draw the priors' center from the Gaussian (uniform for $t_{\rm asc}$) prior centered at the injection parameters with a standard deviation as reported in Table II [19, 39, 40].

We report the results of our optimization in Sec. IV A for a single-stage follow-up and Sec. IV B for multistage follow-up. General discussion is presented in Sec. IV C.

A. Single-stage follow-up

The first follow-up stage uses a coherence time of $T_{\rm coh} = 0.5\,$d, in a similar manner to the O3 all-sky searches [18, 34]. This is a significant increase with respect to the searches' coherence time ($T_{\rm coh} \sim T_{\rm SFT}$), and has been empirically verified to successfully find simulated CW signals given the uncertainties of a typical all-sky search both in real-data searches [18, 33, 40, 41, 45–50] and in broader follow-up studies [22].

Note that this starting value is not given by the coherence-time ladder. This is because previous search stages using shorter coherence times are likely to use different detection statistics that behave differently to the $\mathcal{F}$-statistics; as a result, the first stage is empirically accommodated for the typical uncertainties returned by said stages.

As discussed in Sec. II C 2, we are interested in configurations to find signals with a mismatch compatible with $p_{\rm fa} \lesssim 10^{-2}$ and $p_{\rm fd} \lesssim 10^{-5}$ at the specified sensitivity depths. To do so, for each follow-up configuration, we calculate $\mu^{90\%}$ and the track distance [Eq. (28)] using 100 simulated signals. These results in simulated data are shown in Fig. 4 for both the isolated and binary cases.

To test the robustness of these configurations to deviations from Gaussianity in real noise, the same configurations are evaluated by injecting the same simulated signals in real O3 Advanced LIGO data. The results are shown in Figs. 5. We observe no significant deviation with respect to the Gaussian-noise injection campaign, demonstrating the capabilities of MCMC follow-ups to identify signals in a realistic situation. Note that we still use of the results from Appendix A using Gaussian noise, meaning the false-alarm probability does not correspond to $p_{\rm fa}(\mu, c)$. As we will see in Sec. V, this is not a problem, as real-data searches tend to operate by selecting a specific number of candidates regardless of the actual $p_{\rm fa}$.

Additionally, we compute the distance between the injection point in Fig. 6 for diagnostic purposes. The resulting values are safely below $1/T_{\rm SFT}$.

Our results show that isolated signals can be followed up using configurations involving about an order of magnitude less $\mathcal{F}$-statistic evaluations than previously realized, both in Gaussian and real data. This amounts to about 10^4 $\mathcal{F}$-statistic evaluations at $\mathcal{D} = 50/\sqrt{\rm Hz}$ versus the 1.8×10^5 evaluations used by the standard configuration. For $\mathcal{D} = 30/\sqrt{\rm Hz}$, the number of $\mathcal{F}$-statistic evaluations can be reduced down to about 2×10^3.

A similar reduction is observed for the case of unknown sources in binary systems. Here, however, the sensitivity depth is 20% to 50% lower; a smaller reduction is expected for weaker signals. We also note that the initial prior shifting highly affects the reported performances. The expected deviation of a candidate from the parameters of a true signal is highly dependent on the characteristics of the pipeline producing the candidates in the first place. In this regard, our results are a conservative estimate of the achievable computing cost reduction, and we expect them to be applicable to general all-sky searches. These results are consistent with those reported in [22], where successful follow-ups incur in a growing computing cost as signals become weaker and parameter-space regions become broader.

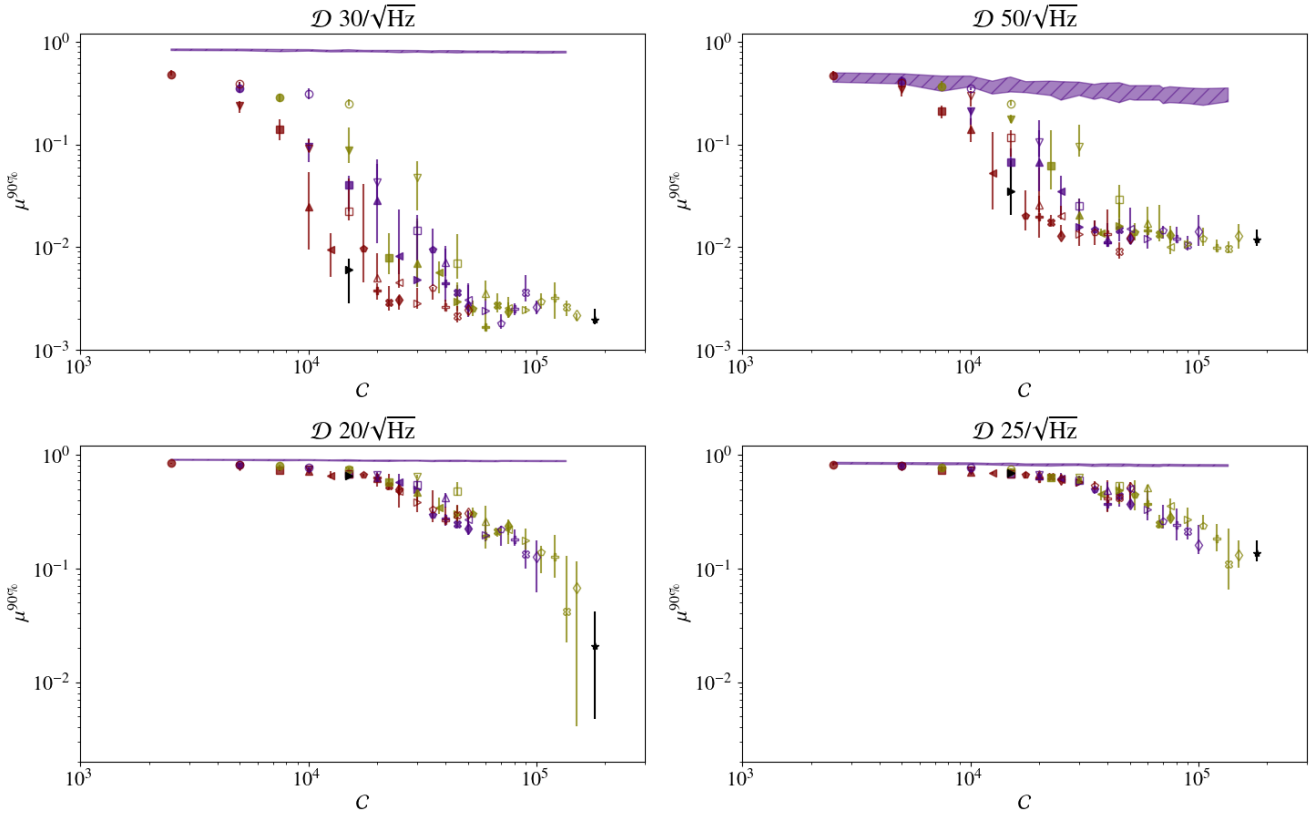

FIG. 4. 90th percentile of the μ distribution as a function of the cost of the MCMC configuration using $T_{\rm coh} = 0.5\,$d. The upper row corresponds to the isolated injection campaign ($T_{\rm obs} = 1$ year), while lower row corresponds to the binary injection campaign ($T_{\rm obs} = 0.5$ year). The shaded regions correspond to μ values such that $10^{-2} \leq p_{\rm fa} \leq 10^{-1}$ and $p_{\rm fd} = 10^{-5}$ at the specified sensitivity depth (see Sec. II C 2 and Appendix A). Each configuration's performance is identified by a different marker. The starred black marker represents the performance of the standard configuration used in previous searches (see text). The black triangular marker denotes the configuration $N_{\rm w} = 50, N_{\rm t} = 1, N_{\rm tot} = 300$ that we identify as a good general choice. Filled (unfilled) markers correspond to $N_w = 50\,(100)$. Red, purple, and green markers correspond to $N_t = 1$, 2, 3, respectively. The shape of the marker denotes different values of $N_{\rm tot}$. Uncertainties are reported at 68% confidence level using bootstrapping. Any configuration below the shaded region is appropriate to successfully follow up a signal at the specified depth level.

	$\mathcal{N}_*$	$T^{(0)}_{\rm coh}$	$T^{(1)}_{\rm coh}$	$T^{(2)}_{\rm coh}$	$T^{(3)}_{\rm coh}$
Isolated	10^4	0.5	5	50	–
Binary	10^3	0.5	1.6	5	15

TABLE III. Coherence times in days for the follow-up of isolated and binary signals calculated with Eq. (26).

B. multistage follow-up

We now test the performance of applying a coherence-time ladder to the results of Sec. IV A. To do so, we select $\mathcal{N}_* = 10^3, 10^4$ for the binary and isolated cases, respectively, and compute the coherence ladder starting from $T_{\rm coh} = 0.5$ day using Eq. (26). The resulting ladders are listed in Table III. As discussed in the introduction, high coherence times may fail to recover a CW signal due to unmodeled effects. To account for these, we limit the maximum coherence time to 50 days in the isolated case to account for possible pulsar glitches [6, 91] and to 15 days in the binary case to account for spin wandering [7, 65, 66].

We compute $\mu^{90\%}$ and $d^{90\%}$ for the different stages as in Sec. IV A. Each stage takes as input the loudest candidate from the previous one. Prior distributions for stage j are built around the loudest candidate's parameters following Table II, adjusting the priors' width according to $T^{(j-1)}_{\rm coh}$.

The evolution of $\mu^{90\%}$ and $d^{90\%}$ across follow-up stages in real data is shown in Figs. 7 and 8. The results for Gaussian data are shown in Appendix B. In this case, our analysis is limited to compare the behavior of different configurations to that of the standard configuration.

Figure 7 displays a clear trend on the effectiveness of MCMC configurations across the coherence-time ladder: most configurations achieving $\mu^{90\%} \lesssim (0.1 - 0.2)$ in the first follow-up stage, regardless of the signal population, are able to converge to the true signal parameters as the ladder progresses, as shown in Fig. 8. Note, however, that such a convergence is conditioned on the steepness of the ladder.

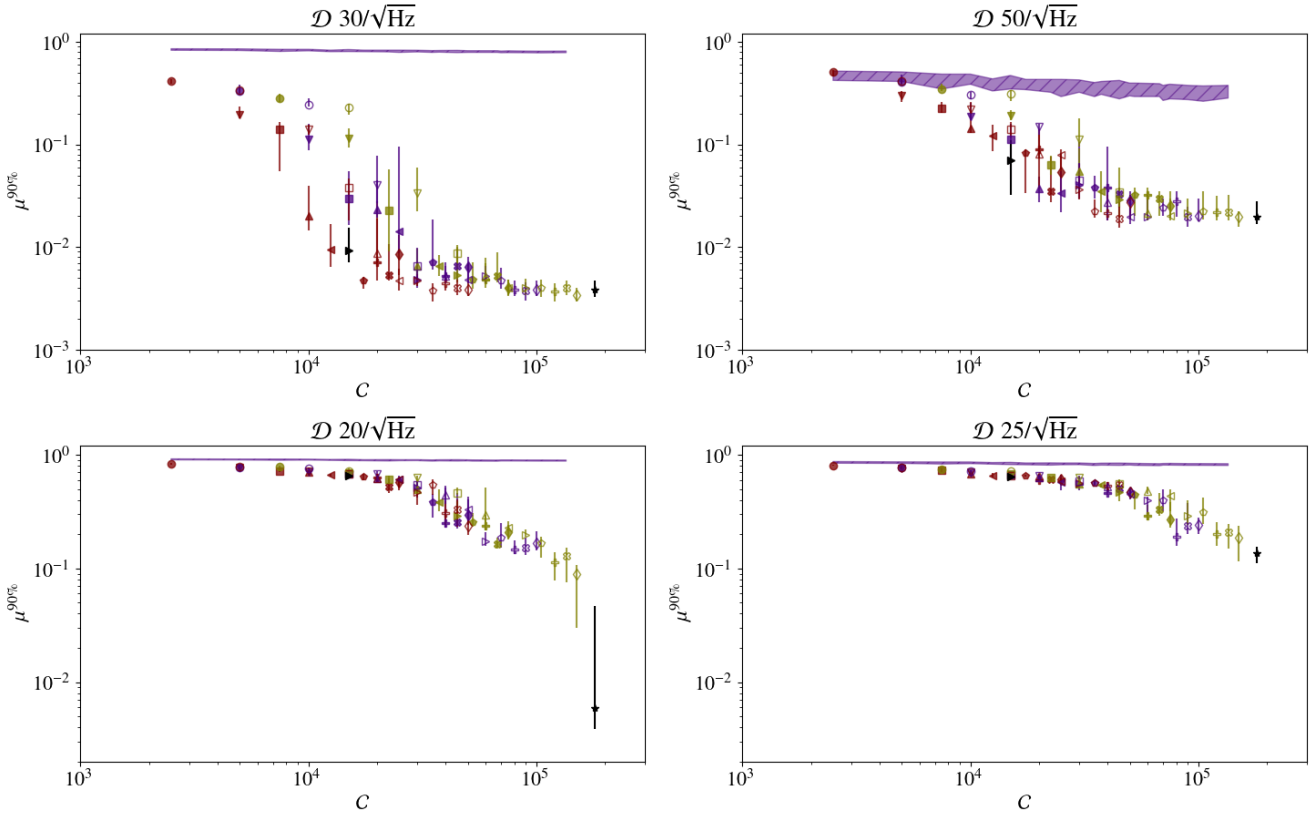

FIG. 5. Same results as in Fig. 4 but using Advanced LIGO O3 and O3a data for isolated (upper row) and binary (lower row) signals, respectively.

C. Discussion

The results presented in Secs. IV A and IV B demonstrate the benefits of optimizing the configuration of the follow-up. In previous works [18, 40, 45, 46, 48, 65, 66], a standard setup was used after verifying proper recovery of the target signal population. Here, we have proved that a successful outcome, with comparable recovery performances, can be obtained at a much lower computing cost.

We identify the configuration $N_{\rm w} = 50, N_{\rm t} = 1, N_{\rm tot} = 300$, corresponding to the black filled right triangle, as a good general default choice for the first follow-up stage, as it reduces the computing cost by a factor 12 and returns an acceptable mismatch level at the probed sensitivity depths for isolated and binary signals in both Gaussian and real data. Cheaper configurations could be selected if one were interested in signals with higher amplitudes ($\mathcal{D} \lesssim 30/\sqrt{\text{Hz}}$). These results are broadly consistent with those discussed in [22] for a different follow-up application.

Figure 7 shows a successful application of our newly proposed coherence time ladder, which is simpler to calculate and tends to return a slightly lower number of steps than the previous proposal. We note, however, that the maximum allowed mismatch for a configuration to be effective across multiple ladder stages is smaller than for a single-stage follow-up. This is especially true for sources in binary systems: due to their higher parameter-space dimensionality, only setups two to five times cheaper than the standard configuration are able to follow the proposed ladder. A more favorable ladder to specific follow-up applications can be constructed by properly tuning the selected value of $\mathcal{N}_*$.

Taken together, these results allow for a variety of computationally efficient strategies. For the case of single-stage follow-ups, the newly proposed configuration will be able to discard noise outliers with the desired false-alarm probability. Upon completion of this first step, slightly more expensive setups (by about a factor 2) can be used to further narrow down the parameters of the candidate before starting the coherence time ladder. Under this scheme, the cost of running a complete ladder on a candidate involves 3 to 5 stages, each of them using 10^3 to, at most, 10^5 $\mathcal{F}$-statistic evaluations.

These results were obtained at a follow-up false-dismissal probability of $p_{\rm fd} = 10^{-5}$ at an amplitude corresponding to the 95% detection probability of a search. In general, however, searches could detect signals at higher sensitivity depths (lower amplitudes), for which the results here presented may be not entirely applicable. We recommend the reader to take the results here presented as well-informed starting point, and to calibrate specific setups for production use by means

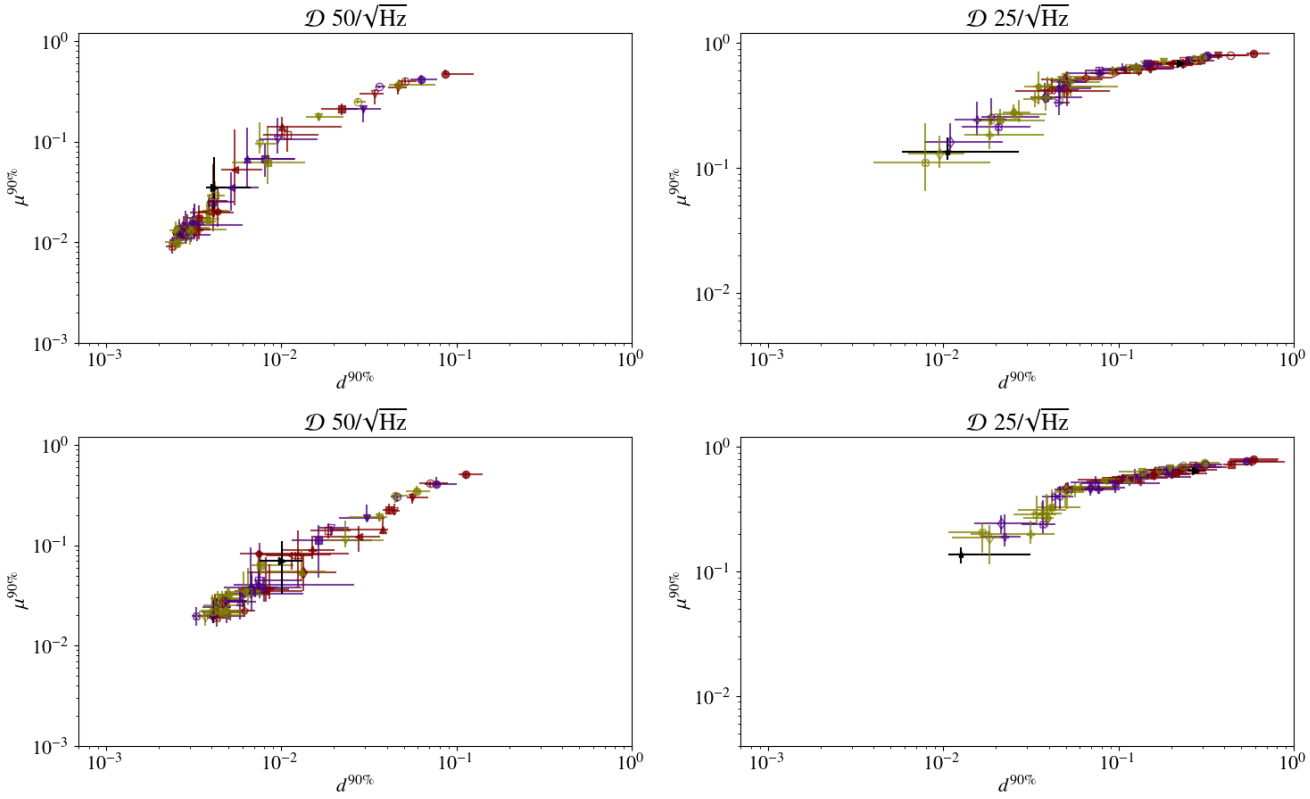

FIG. 6. 90th percentile of the μ distribution as a function of the same percentile of the d distribution for the results presented in Fig. 4. The distance is computed using $T_{\text{SFT}} = 1800$, 1024 s for the isolated (left column) and binary (right column) cases, respectively. Upper panels are computed using Gaussian data; lower panels are computed using real O3 Advanced LIGO data. Uncertainties are reported at the 68% confidence level using bootstrapping. We only show the results for the higher depth values for simplicity.

of a minimal injection campaign to account for any unexpected behaviors of the follow-up caused by the dataset at hand.

V. IMPACT ON ALL-SKY SEARCH SENSITIVITIES

We devote this section to justifying the importance of postprocessing and follow-up strategies when it comes to improving the sensitivity of broad parameter-space searches for CWs. To do so, we run an all-sky search on 40 0.1 Hz representative frequency bands of Advanced LIGO O3 data using the `SkyHough` pipeline [19, 51] with the same configuration as in [18]. These 40 representative frequency bands, namely [66.325, 70.300, 71.275, 73.250, 75.350, 93.200, 99.375, 105.475, 109.600, 119.000, 130.475, 135.650, 142.100, 146.725, 158.225, 168.875, 173.075, 176.150, 181.000, 193.575, 201.325, 215.900, 238.475, 242.875, 247.850, 249.100, 260.200, 262.875, 269.500, 273.650, 280.000, 282.050, 301.000, 304.525, 309.550, 312.900, 322.950, 334.100, 334.900, 336.000] Hz, correspond to a subset of the bands used to estimate the sensitivity of the search.

In each band, an all-sky template bank covering a linear spindown range between -1×10^{-9} Hz/s and 1×10^{-12} Hz/s is analyzed using different detection statistics as described in Sec. IV B 1 of [18]. The result per band is a toplist containing the top 10^5 candidates, which then are clustered using the algorithm presented in [85].

At this point, the search in [18] selected the top 40 clusters in each band and estimated the search sensitivity by means of an injection campaign. The number of selected clusters was significantly higher than previous searches (cf. 1 cluster in O1 isolated [61], 1 cluster in O2 isolated [30], 3 clusters in O2 all-sky binary [33]). In turn, the resulting sensitivity depth increased by about 20% with respect to said searches. Part of this sensitivity increase is due to the duration of the dataset and the improved noise background; part is due to the increase in the number of selected candidates as shown in Fig. 9: An increased number of candidates corresponds to increasing the false-alarm probability of the search, which in turn tends to increase the detected fraction of a given population. This approach, however, runs into diminishing returns as the threshold falls off rather slowly with the number of candidates.

We now evaluate the sensitivity impact of being able to follow up a higher number of CW candidates. To do so, we simulate 200 CW signals at sensitivity depths ranging from $\mathcal{D} = 32/\sqrt{\text{Hz}}$ to $\mathcal{D} = 47/\sqrt{\text{Hz}}$ and interpolate

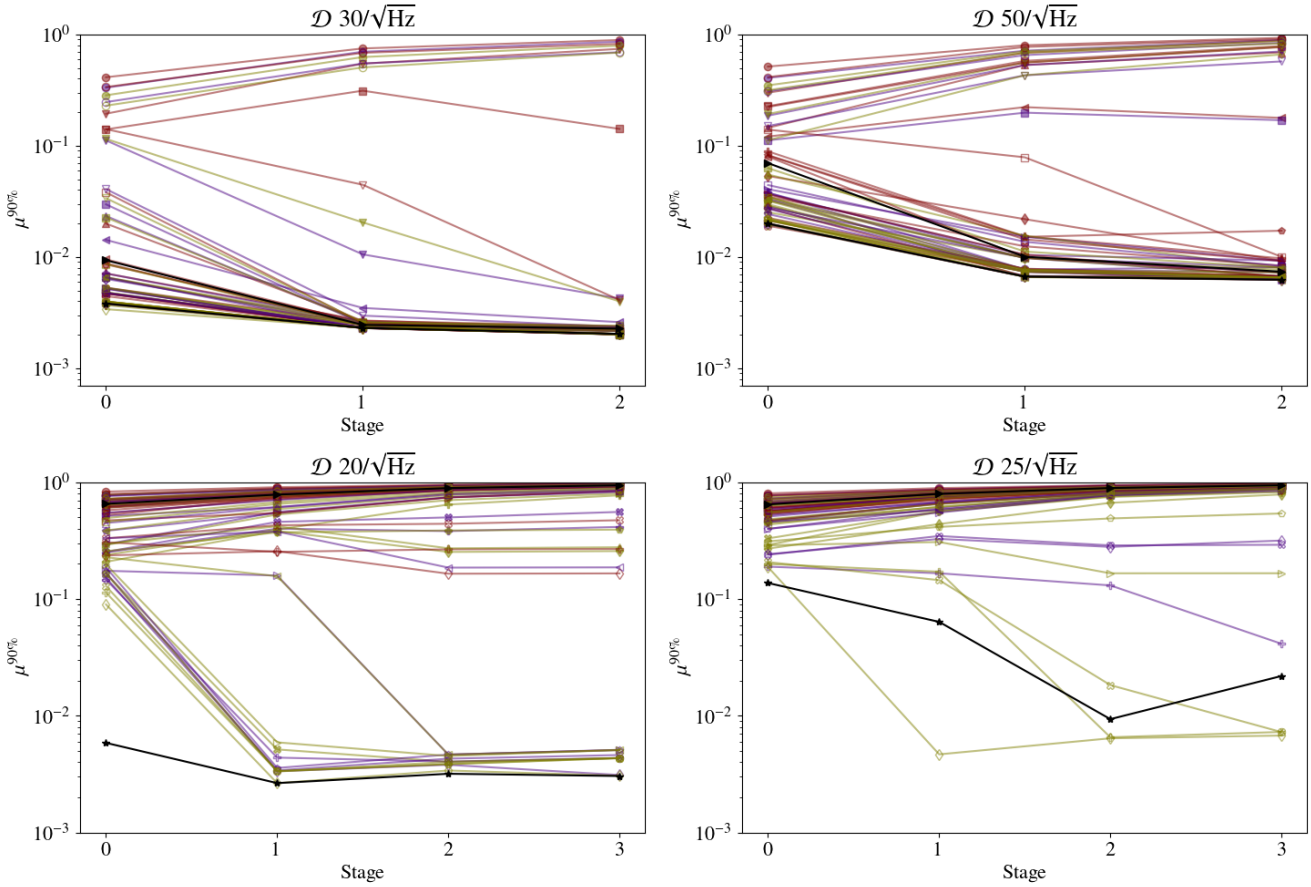

FIG. 7. 90th percentile of the mismatch distribution after each stage of the coherence time ladder (see Table III) on real Advanced LIGO data. The upper row corresponds to the isolated injection campaign. The lower row corresponds to the binary case. Markers are consistent with the legend described in Fig. 4.

the sensitivity depth $\mathcal{D}^{95\%}$ corresponding to a 95% detection probability as explained in Sec. V B 2 of [18]. Uncertainties on $\mathcal{D}^{95\%}$ are computed using the covariance matrix of the fit. We follow a consistent detection criterion: a simulated signal is considered detected if (1) the detection statistic of at least one of the selected clusters is above that of the last selected cluster in the all-sky search and (2) the loudest candidate within said cluster is less than two parameter-space bins away from the injection point. An example of this fitting procedure is shown in Fig. 10.

The resulting sensitivity estimates are summarized in Fig. 11. Increasing the number of selected clusters from 4 to 40 produces a median sensitivity increase of about 10% across the whole frequency band, with a maximum (minimum) increase of 15% (6%). Further increasing from 40 to 400 clusters produces a median increase of about 4%, with a maximum (minimum) increase of 7% (1%).

These results demonstrate that, overall, affordable follow-up strategies imply a better sensitivity for all-sky searches. This improvement, however, is likely to saturate upon reaching the bulk of the noise distribution, as suggested by the threshold behavior in Fig. 9.

VI. CONCLUSION

Blind searches for CW sources are in a privileged position to claim the first detection of a CW signal. This is propitiated by the use of semicoherent methods [5], which allow to evaluate broad parameter-space regions under an affordable computing budget by reducing the strictness with which a signal template is compared to the data [5, 10, 28, 43, 44]. Incidentally, this makes them robust to the stochastic nature of CW sources such as NS stars, which are known to be affected by glitches [6] and spin-wandering [7].

The basic search strategy is to use affordable methods to select interesting candidates which then are further evaluated using more sensitive follow-up methods in a hierarchical scheme. Search sensitivity is thus limited by the number of candidates to be followed up.

In this work, we presented a new framework to evaluate the effectiveness of a follow-up configuration. We apply this framework to `pyfstat` [10, 11], a state-of-the-art MCMC follow-up widely accepted by the CW community [18, 40, 45, 46, 48, 65, 66]. Additionally, we overall simplify the setup of multistage follow-up strategy to ease their applicability to generic CW searches.

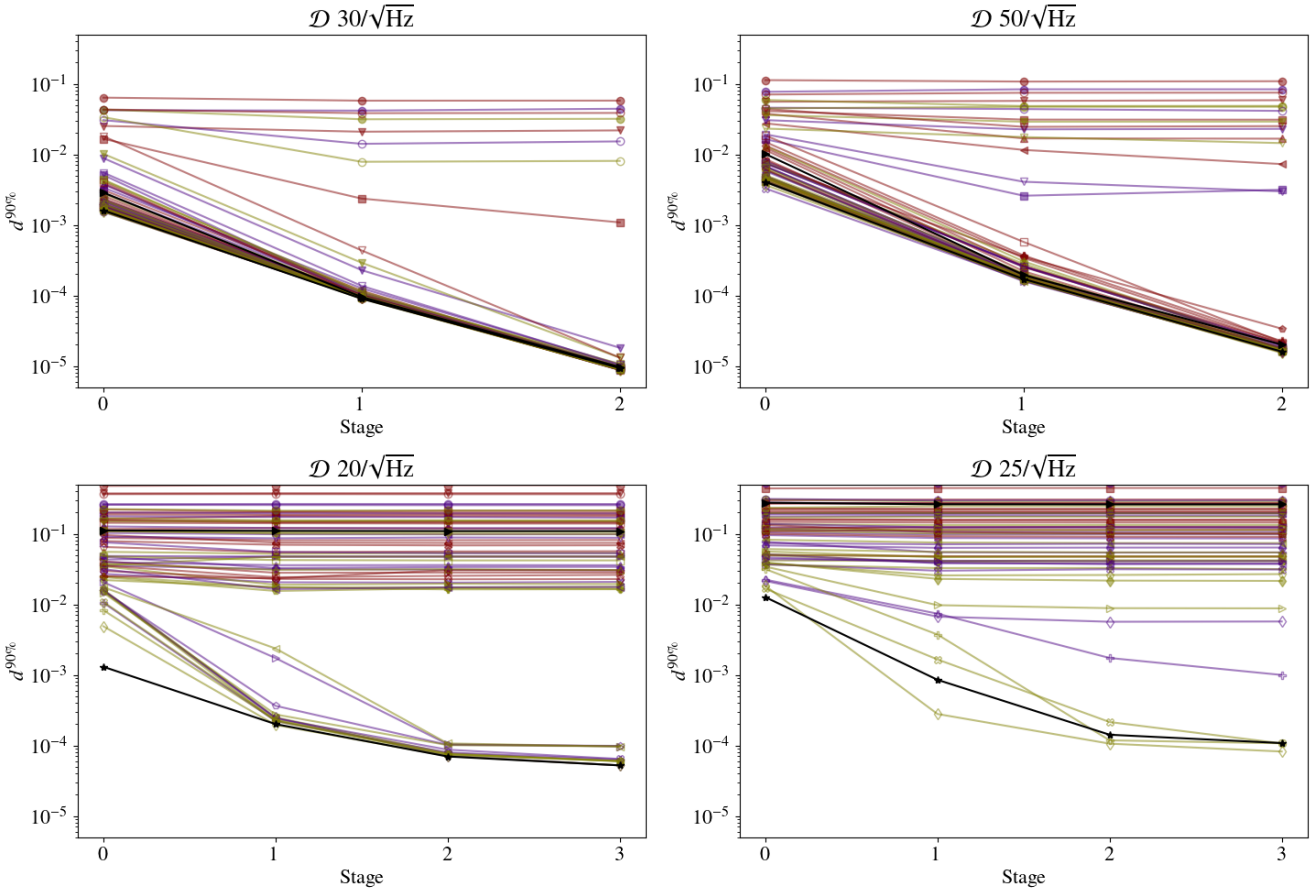

FIG. 8. 90th percentile of the distance distribution after each stage of the coherence time ladder (see Table III) on real Advanced LIGO data. The upper row corresponds to the isolated injection campaign. The lower row corresponds to the binary case. Markers are consistent with the legend described in Fig. 4.

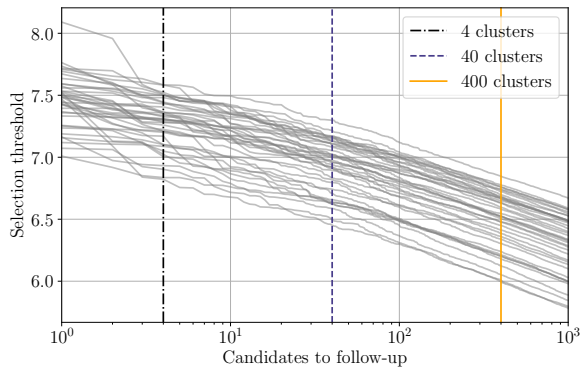

FIG. 9. Statistical threshold corresponding to the last selected cluster for all the analyzed frequency bands as a function of the number of selected candidates. Each gray line represents a different frequency band. The quantity on the vertical axis corresponds to the critical ratio of the weighed normalized power, as described in [18].

We test our proposals using simulated Gaussian and real O3 Advanced LIGO data, and we consider all-sky populations of CW sources both isolated and in binary systems.

We demonstrated significant reduction of the computational cost in the follow-up of generic CW candidates at realistic sensitivity depths on real data. If we focus on relatively strong signals($\mathcal{D} \lesssim 30/\sqrt{\text{Hz}}$), this reduction reaches two orders of magnitude. These findings are consistent with [22], which tackles the problem of broad parameter-space follow-ups. With this, we estimated the sensitivity improvement on an all-sky search by increasing the number of follow-up outliers, which may reach up to a 15% increase in sensitivity depth at a comparable computing cost.

These results will have a positive impact on the resulting sensitivity of all-sky searches in the forthcoming observing runs of the LIGO-Virgo-KAGRA collaboration.

ACKNOWLEDGMENTS

We thank Alessandro De Falco, Pia Astone, Alicia M. Sintes, David Keitel for encouraging and fruitful discussions. We thank Rafel Jaume for computing

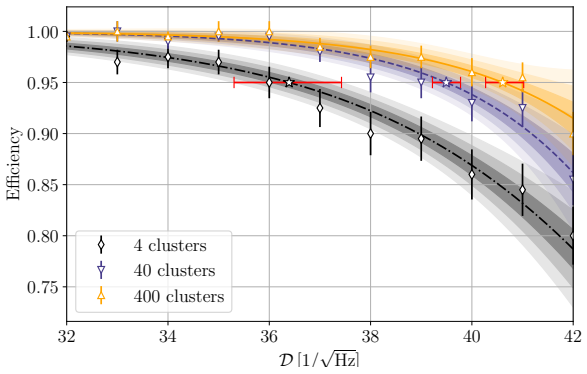

FIG. 10. Example of sensitivity estimation fit. Each set of markers corresponds to a different threshold given by the last selected cluster. The solid, dashed, and dot-dashed lines represent the sigmoid fit. Shaded regions represent the 1, 2, and 3 sigma uncertainties as computed from the covariance matrix. We use one sigma for the 40 and 400 cases, and three sigma for 4 as noise fluctuations make the fits very unstable.

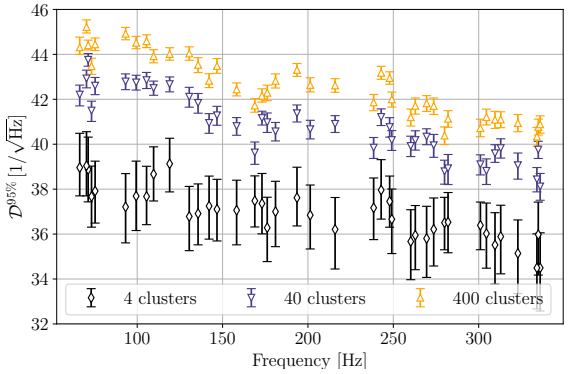

FIG. 11. Sensitivity depths corresponding to a 95% detection efficiency at different frequency bands. Results consistent with the original O3 search [18] are shown as downward triangles. Error bars correspond to 1 sigma fit uncertainties for the cases of 40 and 400 clusters. Due to the instability of the fit, 3 sigma uncertainties were used for the case of 4 clusters.

support and Karl Wette, Luca D'Onofrio, the CW working group of the LIGO-Virgo-KAGRA Collaboration, Pep Covas, Reinhard Prix, Jasper Martins, and the anonymous referee for comments on the manuscript. L.M. thanks the GRAVITY group at the University of the Balearic Islands for their hospitality during the development of this project. R.T. is supported by ERC Starting Grant No. 945155–GWmining, Cariplo Foundation Grant No. 2021-0555, MUR PRIN Grant No. 2022-Z9X4XS, MUR Grant "Progetto Dipartimenti di Eccellenza 2023-2027" (BiCoQ), the ICSC National Research Centre funded by NextGenerationEU, the Universitat de les Illes Balears (UIB); the Spanish Agencia Estatal de Investigación grants PID2022-138626NB-I00, RED2022-134204-E, RED2022-134411-T, funded by MICIU/AEI/10.13039/501100011033, and ERDF/UE; the MICIU with funding from the European Union NextGenerationEU/PRTR (PRTR-C17.I1); the Comunitat Autonòma de les Illes Balears through the Direcció General de Recerca, Innovació I Transformació Digital with funds from the Tourist Stay Tax Law (PDR2020/11 - ITS2017-006), the Conselleria d'Economia, Hisenda i Innovació Grant No. SINCO 2022/18146 (HiTech-IAC3-BIO) cofinanced by the European Union and FEDER Operational Program 2021-2027 of the Balearic Islands. The authors are grateful for computational resources provided by LIGO Laboratory, supported by National Science Foundation Grants No. PHY-0757058 and No. PHY-0823459. This research has made use of data or software obtained from the Gravitational Wave Open Science Center (gwosc.org), a service of the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA. This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation, as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain. KAGRA is supported by Ministry of Education, Culture, Sports, Science and Technology (MEXT), Japan Society for the Promotion of Science (JSPS) in Japan; National Research Foundation (NRF) and Ministry of Science and ICT (MSIT) in Korea; Academia Sinica (AS) and National Science and Technology Council (NSTC) in Taiwan. This document has been assigned document number LIGO-P2400221.

Appendix A: MCMC NOISE DISTRIBUTION

We empirically estimate the noise distribution of different MCMC follow-up configurations. To do so, we run a follow-up using a subset of the MCMC configurations in Table I and priors from Table II (with width $T_{\text{coh}} = T_{\text{SFT}}$) on 150 Gaussian-noise realizations. We select durations of $T_{\text{obs}} = 0.5, 1$ year in order to consider both the binary and isolated cases. The number of segments is 365 and 730, respectively, in order to have $T_{\text{coh}} = 0.5\,\text{d}$.

For each MCMC configuration, we retrieve $2\mathcal{F}_{\text{max}}$ from each of the 150 noise realizations and fit a Gumbel distribution [Eq. (18)] as explained in Sec. II. The

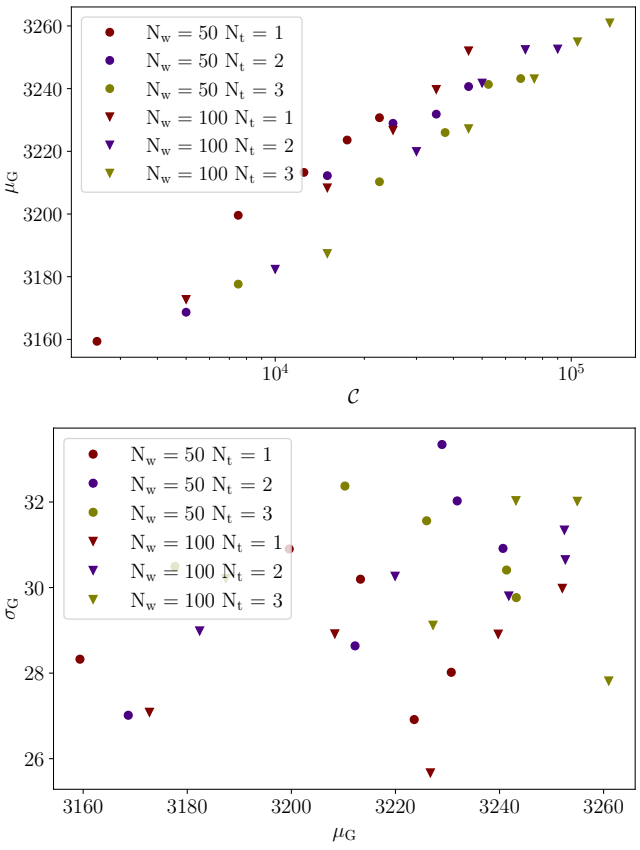

FIG. 12. Upper panel: location parameter of the Gumbel distribution as a function of computing cost. Lower panel: scatter plot of location and scale parameters of a Gumbel distribution. Each pair of parameters is estimated by retrieving the loudest $\mathcal{F}$-statistic of 150 MCMC runs on Gaussian noise and fitting a Gumbel distribution.

resulting location μ_G and scale σ_G parameters for $T_{obs} = 1$ year are shown in Fig. 12. Similar results are observed for $T_{obs} = 0.5$ year.

We find a strong correlation between the location parameter μ_G and (the logarithm) of the cost of the configuration. This is expected, as, in this case, μ_G grows with the logarithm of the size of the template bank [69], which is proportional to the cost of a configuration. The scale parameter fluctuates by about 10% around an average value of ~ 30.

The noise distribution of the $\mathcal{F}$-statistic in Gaussian noise depends only on N_{seg} regardless of the number of detectors, their PSD, and the observing time [35–37]. In particular, the results generated in this appendix can be used for both the isolated and binary injection campaigns shown in Figs. 4 and 5. We note, however, that the behavior of outliers on real data is difficult to predict as it strongly depends on their specific cause.

Appendix B: MULTISTAGE FOLLOW-UP WITH SIMULATED DATA

As mentioned in Sec. IV, we show in Fig. 13 and Fig. 14 the mismatch and distance results for a multistage follow-up on simulated Gaussian data.

[1] K. Riles, Searches for continuous-wave gravitational radiation, Living Rev. Rel. **26**, 3 (2023), arXiv:2206.06447 [astro-ph.HE].

[2] B. T. Reed, A. Deibel, and C. J. Horowitz, Modeling the Galactic Neutron Star Population for Use in Continuous Gravitational-wave Searches, Astrophys. J. **921**, 89 (2021), arXiv:2104.00771 [astro-ph.HE].

[3] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, The Australia Telescope National Facility Pulsar Catalogue, Astron. J. **129**, 1993 (2005), arXiv:astro-ph/0412641 [astro-ph].

[4] K. Wette, Searches for continuous gravitational waves from neutron stars: A twenty-year retrospective, Astropart. Phys. **153**, 102880 (2023), arXiv:2305.07106 [gr-qc].

[5] R. Tenorio, D. Keitel, and A. M. Sintes, Search Methods for Continuous Gravitational-Wave Signals from Unknown Sources in the Advanced-Detector Era, Universe **7**, 474 (2021), arXiv:2111.12575 [gr-qc].

[6] G. Ashton, R. Prix, and D. I. Jones, Statistical characterization of pulsar glitches and their potential impact on searches for continuous gravitational waves, Phys. Rev. D **96**, 063004 (2017), arXiv:1704.00742 [gr-qc].

[7] A. Mukherjee, C. Messenger, and K. Riles, Accretion-induced spin-wandering effects on the neutron star in Scorpius X-1: Implications for continuous gravitational wave searches, Phys. Rev. D **97**, 043016 (2018), arXiv:1710.06185 [gr-qc].

[8] W. D. Vousden, W. M. Farr, and I. Mandel, Dynamic temperature selection for parallel tempering in Markov chain Monte Carlo simulations, Mon. Not. R. Astron. Soc. **455**, 1919 (2016), arXiv:1501.05823 [astro-ph.IM].

[9] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, emcee: The MCMC Hammer, Publ. Astron. Soc. Pac. **125**, 306 (2013), arXiv:1202.3665 [astro-ph.IM].

[10] G. Ashton and R. Prix, Hierarchical multistage MCMC follow-up of continuous gravitational wave candidates, Phys. Rev. D **97**, 103020 (2018), arXiv:1802.05450 [astro-ph.IM].

[11] D. Keitel, R. Tenorio, G. Ashton, and R. Prix, PyFstat: a Python package for continuous gravitational-wave data analysis, J. Open Source Softw. **6**, 3000 (2021), arXiv:2101.10915 [gr-qc].

[12] M. Tse *et al.*, Quantum-Enhanced Advanced LIGO Detectors in the Era of Gravitational-Wave Astronomy,

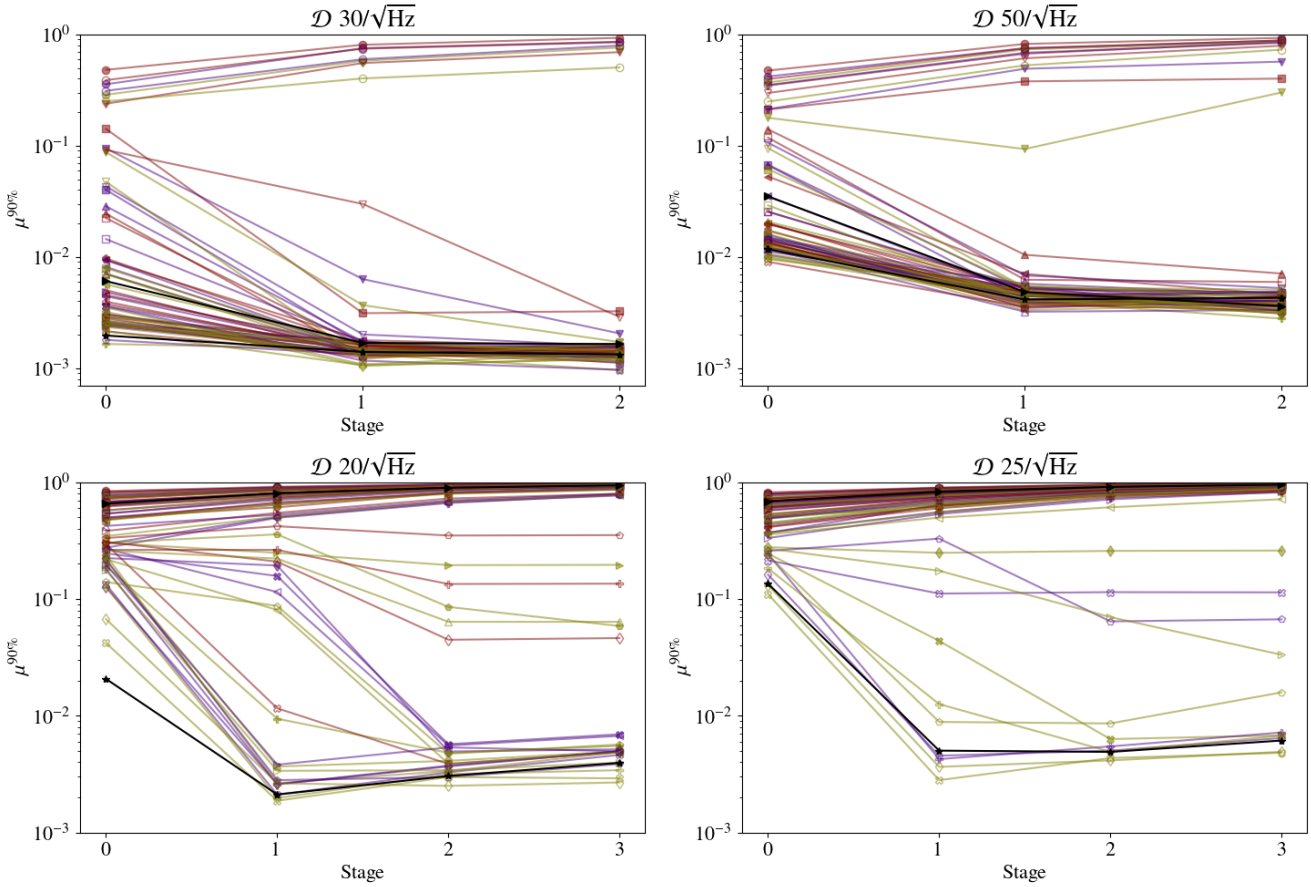

FIG. 13. Same caption as Fig. 7, but using Gaussian data.

Phys. Rev. Lett. **123**, 231107 (2019).

[13] F. Acernese *et al.* (Virgo), Increasing the Astrophysical Reach of the Advanced Virgo Detector via the Application of Squeezed Vacuum States of Light, Phys. Rev. Lett. **123**, 231108 (2019).

[14] A. Buikema *et al.*, Sensitivity and performance of the Advanced LIGO detectors in the third observing run, Phys. Rev. D **102**, 062003 (2020), arXiv:2008.01301 [astro-ph.IM].

[15] D. Davis *et al.* (LIGO), LIGO detector characterization in the second and third observing runs, Class. Quant. Grav. **38**, 135014 (2021), arXiv:2101.11673 [astro-ph.IM].

[16] F. Acernese *et al.* (Virgo), Virgo detector characterization and data quality: results from the O3 run, Class. Quant. Grav. **40**, 185006 (2023), arXiv:2210.15633 [gr-qc].

[17] R. Abbott *et al.* (KAGRA, Virgo, LIGO Scientific), Open Data from the Third Observing Run of LIGO, Virgo, KAGRA, and GEO, Astrophys. J. Suppl. **267**, 29 (2023), arXiv:2302.03676 [gr-qc].

[18] R. Abbott *et al.* (KAGRA, LIGO Scientific, Virgo), All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO and Advanced Virgo O3 data, Phys. Rev. D **106**, 102008 (2022), arXiv:2201.00697 [gr-qc].

[19] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, and C. Palomba, The Hough transform search for continuous gravitational waves, Phys. Rev. D **70**, 082001 (2004), arXiv:gr-qc/0407001.

[20] P. Astone, A. Colla, S. D'Antonio, S. Frasca, and C. Palomba, Method for all-sky searches of continuous gravitational wave signals using the frequency-hough transform, Phys. Rev. D **90**, 042002 (2014), arXiv:1407.8333 [astro-ph.IM].

[21] P. B. Covas and A. M. Sintes, New method to search for continuous gravitational waves from unknown neutron stars in binary systems, Phys. Rev. D **99**, 124019 (2019), arXiv:1904.04873 [astro-ph.IM].

[22] P. B. Covas, R. Prix, and J. Martins, New framework to follow up candidates from continuous gravitational-wave searches, Phys. Rev. D **110**, 024053 (2024), arXiv:2404.18608 [gr-qc].

[23] P. R. Brady and T. Creighton, Searching for periodic sources with LIGO. 2. Hierarchical searches, Phys. Rev. D **61**, 082001 (2000), arXiv:gr-qc/9812014.

[24] C. Cutler, I. Gholami, and B. Krishnan, Improved stack-slide searches for gravitational-wave pulsars, Phys. Rev. D **72**, 042004 (2005), arXiv:gr-qc/0505082.

[25] R. Prix and M. Shaltev, Search for Continuous Gravitational Waves: Optimal StackSlide method at fixed computing cost, Phys. Rev. D **85**, 084010 (2012), arXiv:1201.4321 [gr-qc].

[26] M. Shaltev and R. Prix, Fully coherent follow-up of continuous gravitational-wave candidates, Phys. Rev. D **87**, 084057 (2013), arXiv:1303.2471 [gr-qc].

[27] M. Shaltev, P. Leaci, M. A. Papa, and R. Prix, Fully coherent follow-up of continuous gravitational-wave candidates: an application to Einstein@Home results,

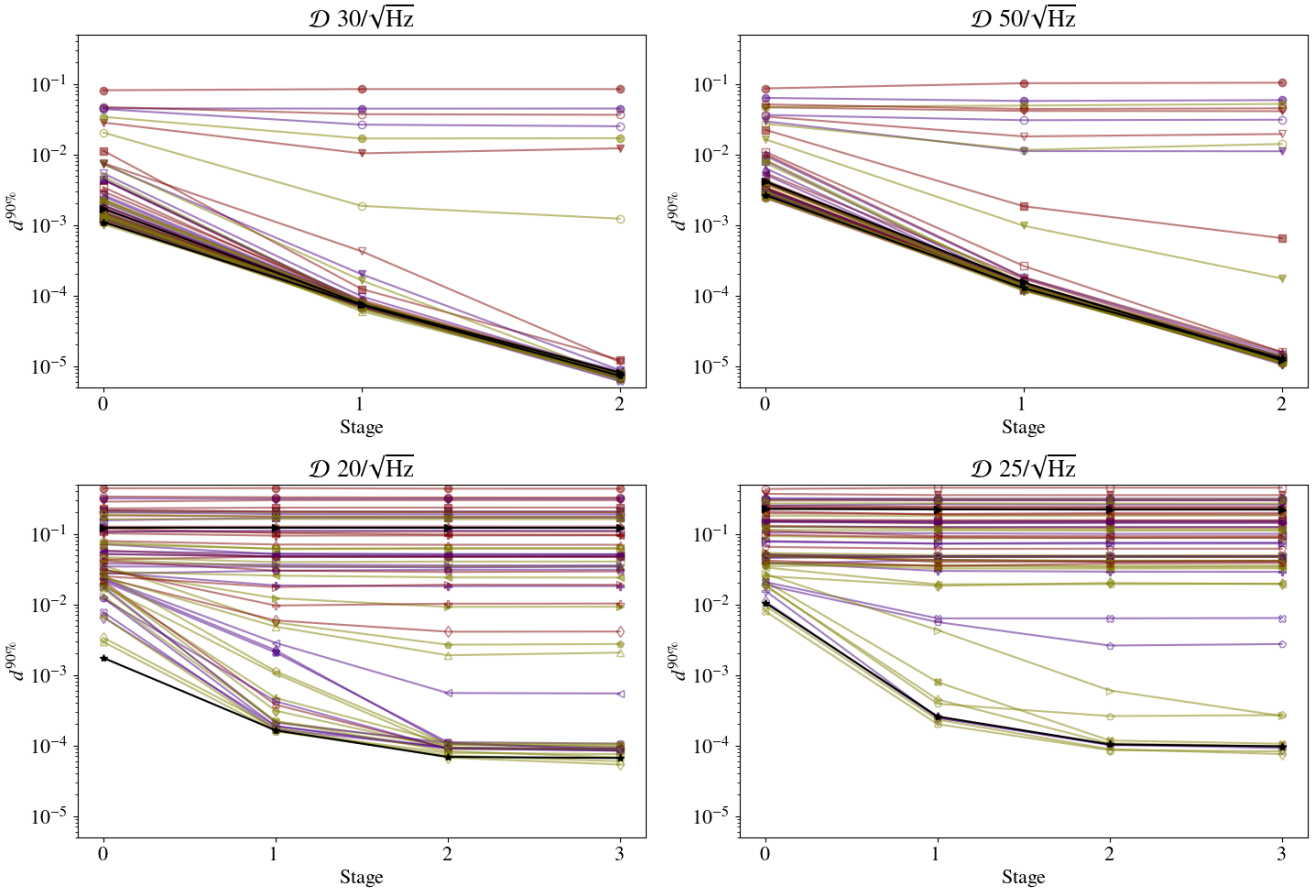

FIG. 14. Same caption as Fig. 8, but using Gaussian data.

Phys. Rev. D **89**, 124030 (2014), arXiv:1405.1922 [gr-qc].
[28] M. Shaltev, Optimizing the StackSlide setup and data selection for continuous-gravitational-wave searches in realistic detector data, Phys. Rev. D **93**, 044058 (2016), arXiv:1510.06427 [gr-qc].
[29] M. A. Papa *et al.*, Hierarchical follow-up of subthreshold candidates of an all-sky Einstein@Home search for continuous gravitational waves on LIGO sixth science run data, Phys. Rev. D **94**, 122006 (2016), arXiv:1608.08928 [astro-ph.IM].
[30] B. P. Abbott *et al.* (LIGO Scientific, Virgo), All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO O2 data, Phys. Rev. D **100**, 024004 (2019), arXiv:1903.01901 [astro-ph.HE].
[31] B. Steltner, M. A. Papa, H. B. Eggenstein, B. Allen, V. Dergachev, R. Prix, B. Machenschalk, S. Walsh, S. J. Zhu, and S. Kwang, Einstein@Home All-sky Search for Continuous Gravitational Waves in LIGO O2 Public Data, Astrophys. J. **909**, 79 (2021), arXiv:2009.12260 [astro-ph.HE].
[32] B. Steltner, M. A. Papa, H. B. Eggenstein, R. Prix, M. Bensch, B. Allen, and B. Machenschalk, Deep Einstein@Home All-sky Search for Continuous Gravitational Waves in LIGO O3 Public Data, Astrophys. J. **952**, 55 (2023), arXiv:2303.04109 [gr-qc].
[33] P. B. Covas and A. M. Sintes, First all-sky search for continuous gravitational-wave signals from unknown neutron stars in binary systems using Advanced LIGO data, Phys. Rev. Lett. **124**, 191102 (2020), arXiv:2001.08411 [gr-qc].
[34] R. Tenorio, D. Keitel, and A. M. Sintes, Application of a hierarchical MCMC follow-up to Advanced LIGO continuous gravitational-wave candidates, Phys. Rev. D **104**, 084012 (2021), arXiv:2105.13860 [gr-qc].
[35] P. Jaranowski, A. Krolak, and B. F. Schutz, Data analysis of gravitational - wave signals from spinning neutron stars. 1. The Signal and its detection, Phys. Rev. D **58**, 063001 (1998), arXiv:gr-qc/9804014.
[36] C. Cutler and B. F. Schutz, The Generalized F-statistic: Multiple detectors and multiple GW pulsars, Phys. Rev. D **72**, 063006 (2005), arXiv:gr-qc/0504011.
[37] R. Prix and B. Krishnan, Targeted search for continuous gravitational waves: Bayesian versus maximum-likelihood statistics, Class. Quant. Grav. **26**, 204013 (2009), arXiv:0907.2569 [gr-qc].
[38] J. T. Whelan, R. Prix, C. J. Cutler, and J. L. Willis, New Coordinates for the Amplitude Parameter Space of Continuous Gravitational Waves, Class. Quant. Grav. **31**, 065002 (2014), arXiv:1311.0065 [gr-qc].
[39] P. Leaci and R. Prix, Directed searches for continuous gravitational waves from binary systems: parameter-space metrics and optimal Scorpius X-1 sensitivity, Phys. Rev. D **91**, 102003 (2015), arXiv:1502.00914 [gr-qc].
[40] R. Abbott *et al.* (LIGO Scientific, Virgo), All-sky search in early O3 LIGO data for continuous gravitational-wave signals from unknown neutron stars in binary systems, Phys. Rev. D **103**, 064017 (2021), [Erratum: Phys.Rev.D

108, 069901 (2023)], arXiv:2012.12128 [gr-qc].
[41] P. B. Covas, M. A. Papa, R. Prix, and B. J. Owen, Constraints on r-modes and Mountains on Millisecond Neutron Stars in Binary Systems, Astrophys. J. Lett. **929**, L19 (2022), arXiv:2203.01773 [gr-qc].
[42] R. Prix, S. Giampanis, and C. Messenger, Search method for long-duration gravitational-wave transients from neutron stars, Phys. Rev. D **84**, 023007 (2011), arXiv:1104.1704 [gr-qc].
[43] V. Dergachev and M. A. Papa, Sensitivity improvements in the search for periodic gravitational waves using O1 LIGO data, Phys. Rev. Lett. **123**, 101101 (2019), arXiv:1902.05530 [gr-qc].
[44] A. J. K. Chua, One-stop function for gravitational-wave detection, identification, and inference, Phys. Rev. D **106**, 104051 (2022), arXiv:2205.08702 [gr-qc].
[45] R. Abbott *et al.* (KAGRA, Virgo, LIGO Scientific), All-sky search for continuous gravitational waves from isolated neutron stars in the early O3 LIGO data, Phys. Rev. D **104**, 082004 (2021), arXiv:2107.00600 [gr-qc].
[46] R. Abbott *et al.* (KAGRA, LIGO Scientific, Virgo), Search for continuous gravitational wave emission from the Milky Way center in O3 LIGO-Virgo data, Phys. Rev. D **106**, 042003 (2022), arXiv:2204.04523 [astro-ph.HE].
[47] R. Abbott *et al.* (LIGO Scientific, Virgo, KAGRA), Search for continuous gravitational waves from 20 accreting millisecond x-ray pulsars in O3 LIGO data, Phys. Rev. D **105**, 022002 (2022), arXiv:2109.09255 [astro-ph.HE].
[48] R. Abbott *et al.* (LIGO Scientific, KAGRA, Virgo), Model-based Cross-correlation Search for Gravitational Waves from the Low-mass X-Ray Binary Scorpius X-1 in LIGO O3 Data, Astrophys. J. Lett. **941**, L30 (2022), arXiv:2209.02863 [astro-ph.HE].
[49] J. T. Whelan *et al.*, Search for Gravitational Waves from Scorpius X-1 in LIGO O3 Data with Corrected Orbital Ephemeris, Astrophys. J. **949**, 117 (2023), arXiv:2302.10338 [astro-ph.HE].
[50] A. Tripathee and K. Riles, Probing more deeply in an all-sky search for continuous gravitational waves in the LIGO O3 data set, Phys. Rev. D **109**, 043049 (2024), arXiv:2311.04985 [gr-qc].
[51] LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration, LVK Algorithm Library - LALSuite, Free software (GPL) (2018).
[52] K. Wette, SWIGLAL: Python and Octave interfaces to the LALSuite gravitational-wave data analysis libraries, SoftwareX **12**, 100634 (2020).
[53] K. Wette, Estimating the sensitivity of wide-parameter-space searches for gravitational-wave pulsars, Phys. Rev. D **85**, 042003 (2012), arXiv:1111.5650 [gr-qc].
[54] B. Behnke, M. A. Papa, and R. Prix, Postprocessing methods used in the search for continuous gravitational-wave signals from the Galactic Center, Phys. Rev. D **91**, 064007 (2015), arXiv:1410.5997 [gr-qc].
[55] C. Dreissigacker, R. Prix, and K. Wette, Fast and Accurate Sensitivity Estimation for Continuous-Gravitational-Wave Searches, Phys. Rev. D **98**, 084058 (2018), arXiv:1808.02459 [gr-qc].
[56] J. Ming, M. A. Papa, H.-B. Eggenstein, B. Beheshtipour, B. Machenschalk, R. Prix, B. Allen, and M. Bensch, Deep Einstein@Home search for Continuous Gravitational Waves from the Central Compact Objects in the Supernova Remnants Vela Jr. and G347.3-0.5 using LIGO public data, arXiv e-prints (2024), arXiv:2408.14573 [gr-qc].
[57] R. Tenorio and L. Mirasola, cows3: continuous-wave search sensitivity simulator, `https://github.com/Rodrigo-Tenorio/continuous-wave-search-sensitivity-simulator` (2024).
[58] P. Astone, S. Frasca, and C. Palomba, The short FFT database and the peak map for the hierarchical search of periodic sources, Class. Quant. Grav. **22**, S1197 (2005).
[59] R. Prix, F-statistic bias due to noise-estimator, `https://dcc.ligo.org/LIGO-T1100551`.
[60] K. Wette, R. Prix, D. Keitel, M. Pitkin, C. Dreissigacker, J. T. Whelan, and P. Leaci, OctApps: a library of Octave functions for continuous gravitational-wave data analysis, J. Open Source Softw. **3**, 707 (2018), arXiv:1806.07442 [astro-ph.IM].
[61] B. P. Abbott *et al.* (LIGO Scientific, Virgo), All-sky Search for Periodic Gravitational Waves in the O1 LIGO Data, Phys. Rev. D **96**, 062002 (2017), arXiv:1707.02667 [gr-qc].
[62] B. P. Abbott *et al.* (LIGO Scientific, Virgo), Full Band All-sky Search for Periodic Gravitational Waves in the O1 LIGO Data, Phys. Rev. D **97**, 102003 (2018), arXiv:1802.05241 [gr-qc].
[63] P. B. Covas, M. A. Papa, and R. Prix, Search for continuous gravitational waves from unknown neutron stars in binary systems with long orbital periods in O3 data, arXiv e-prints (2024), arXiv:2409.16196 [gr-qc].
[64] J. Aasi *et al.* (LIGO Scientific), Advanced LIGO, Class. Quant. Grav. **32**, 074001 (2015), arXiv:1411.4547 [gr-qc].
[65] R. Abbott *et al.* (LIGO Scientific, Virgo, KAGRA), Search for continuous gravitational waves from 20 accreting millisecond x-ray pulsars in O3 LIGO data, Phys. Rev. D **105**, 022002 (2022), arXiv:2109.09255 [astro-ph.HE].
[66] J. T. Whelan *et al.*, Search for Gravitational Waves from Scorpius X-1 in LIGO O3 Data with Corrected Orbital Ephemeris, Astrophys. J. **949**, 117 (2023), arXiv:2302.10338 [astro-ph.HE].
[67] D. Keitel, G. Woan, M. Pitkin, C. Schumacher, B. Pearlstone, K. Riles, A. G. Lyne, J. Palfreyman, B. Stappers, and P. Weltevrede, First search for long-duration transient gravitational waves after glitches in the Vela and Crab pulsars, Phys. Rev. D **100**, 064058 (2019), arXiv:1907.04717 [gr-qc].
[68] K. Wette, L. Dunn, P. Clearwater, and A. Melatos, Deep exploration for continuous gravitational waves at 171–172 Hz in LIGO second observing run data, Phys. Rev. D **103**, 083020 (2021), arXiv:2103.12976 [gr-qc].
[69] R. Tenorio, L. M. Modafferi, D. Keitel, and A. M. Sintes, Empirically estimating the distribution of the loudest candidate from a gravitational-wave search, Phys. Rev. D **105**, 044029 (2022), arXiv:2111.12032 [gr-qc].
[70] R. Abbott *et al.* (LIGO Scientific, KAGRA, Virgo), Narrowband Searches for Continuous and Long-duration Transient Gravitational Waves from Known Pulsars in the LIGO-Virgo Third Observing Run, Astrophys. J. **932**, 133 (2022), arXiv:2112.10990 [gr-qc].
[71] R. Jaume, R. Tenorio, and A. M. Sintes, Assessing the Similarity of Continuous Gravitational-Wave Signals to Narrow Instrumental Artifacts, Universe **10**, 121 (2024), arXiv:2403.03027 [gr-qc].
[72] M. Isi, S. Mastrogiovanni, M. Pitkin, and O. J. Piccinni, Establishing the significance of continuous gravitational-wave detections from known pulsars, Phys. Rev. D **102**,

123027 (2020), arXiv:2010.12612 [gr-qc].
[73] P. Virtanen *et al.*, SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, Nature Methods **17**, 261 (2020).
[74] K. Wette, Empirically extending the range of validity of parameter-space metrics for all-sky searches for gravitational-wave pulsars, Phys. Rev. D **94**, 122002 (2016), arXiv:1607.00241 [gr-qc].
[75] N. Chopin and O. Papaspiliopoulos, *An introduction to Sequential Monte Carlo*, 1st ed., Springer Series in Statistics (Springer, Cham, Switzerland, 2020).
[76] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz, Searching for periodic sources with LIGO, Phys. Rev. D **57**, 2101 (1998), arXiv:gr-qc/9702050.
[77] P. Jaranowski and A. Krolak, Data analysis of gravitational wave signals from spinning neutron stars. 2. Accuracy of estimation of parameters, Phys. Rev. D **59**, 063003 (1999), arXiv:gr-qc/9809046.
[78] R. Prix, Search for continuous gravitational waves: Metric of the multidetector $\mathcal{F}$-statistic, Phys. Rev. D **75**, 023004 (2007), [Erratum: Physical Review D 75, 069901(E) (2007)].
[79] H. J. Pletsch, Parameter-space correlations of the optimal statistic for continuous gravitational-wave detection, Phys. Rev. D **78**, 102005 (2008), arXiv:0807.1324 [gr-qc].
[80] H. J. Pletsch, Parameter-space metric of semicoherent searches for continuous gravitational waves, Phys. Rev. D **82**, 042002 (2010), arXiv:1005.0395 [gr-qc].
[81] K. Wette and R. Prix, Flat parameter-space metric for all-sky searches for gravitational-wave pulsars, Phys. Rev. D **88**, 123005 (2013), arXiv:1310.5587 [gr-qc].
[82] K. Wette, Lattice template placement for coherent all-sky searches for gravitational-wave pulsars, Phys. Rev. D **90**, 122010 (2014), arXiv:1410.6882 [gr-qc].
[83] K. Wette, Parameter-space metric for all-sky semicoherent searches for gravitational-wave pulsars, Phys. Rev. D **92**, 082003 (2015), arXiv:1508.02372 [gr-qc].
[84] M. T. M. Emmerich and A. H. Deutz, A tutorial on multiobjective optimization: fundamentals and evolutionary methods, Natural Computing **17**, 585 (2018).
[85] R. Tenorio, D. Keitel, and A. M. Sintes, Time-frequency track distance for comparing continuous gravitational wave signals, Phys. Rev. D **103**, 064053 (2021), arXiv:2012.05752 [gr-qc].
[86] E. Goetz and K. Riles, Segments used for creating standard SFTs in O3 data, `https://dcc.ligo.org/LIGO-T2300068/public`, LIGO-T2300068.
[87] J. Zweizig and K. Riles, Information on self-gating of $h(t)$ used in O3 continuous-wave and stochastic searches, `https://dcc.ligo.org/LIGO-T2000384/public`, LIGO-T2000384.
[88] B. Allen, E. Goetz, D. Keitel, M. Landry, G. Mendell, R. Prix, K. Riles, and K. Wette, SFT Data Format Version 2–3 Specification, `https://dcc.ligo.org/LIGO-T040164`, T040164.
[89] R. Prix, The F-statistic and its implementation in ComputeFstatistic_v2, `https://dcc.ligo.org/LIGO-T0900149/public` (2019).
[90] E. T. Jaynes, *Probability Theory: The Logic of Science*, edited by G. L. Bretthorst and Cambridge University Press (Cambridge University Press, 2003).
[91] C. M. Espinoza, A. G. Lyne, B. W. Stappers, and M. Kramer, A study of 315 glitches in the rotation of 102 pulsars, MNRAS **414**, 1679 (2011), arXiv:1102.1743 [astro-ph.HE].