

Validating continuous gravitational-wave candidates from a semicoherent search using Doppler modulation and an effective point spread function

Dana Jones^{1,2,*}, Ling Sun², Julian Carlin^{3,4}, Liam Dunn^{3,4}, Meg Millhouse^{3,4}, Hannah Middleton^{3,4,5,6}, Patrick Meyers^{3,4,7}, Patrick Clearwater^{3,4}, Deeksha Beniwal⁸, Lucy Strang^{3,4}, Andrés Vargas^{3,4} and Andrew Melatos^{3,4}

¹*College of Arts and Sciences, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA*

²*OzGrav-ANU, Centre for Gravitational Astrophysics, College of Science, The Australian National University, Australian Capital Territory 2601, Australia*

³*School of Physics, University of Melbourne, Parkville, Victoria 3010, Australia*


⁴*OzGrav, University of Melbourne, Parkville, Victoria 3010, Australia*

⁵*Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn, Victoria 3122, Australia*

⁶*School of Physics and Astronomy and Institute for Gravitational Wave Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, United Kingdom*

⁷*Theoretical Astrophysics Group, California Institute of Technology, Pasadena, California 91125, USA*

⁸*OzGrav, University of Adelaide, Adelaide, South Australia 5005, Australia*

 (Received 28 March 2022; revised 19 October 2022; accepted 7 November 2022; published 13 December 2022)

Following up large numbers of candidates in continuous gravitational-wave searches presents a challenge, particularly in regard to computational power and the time required to manually scrutinize each of the candidates. It is important to design and test good follow-up procedures that are safe (i.e., minimize false dismissals) and computationally efficient across many search configurations. We investigate two follow-up procedures, or “vetoes,” both of which exploit the Doppler modulation predicted in astrophysical signals. In particular, we introduce the concept of using an effective point spread function as part of our veto criteria. We take advantage of a well-established semicoherent search algorithm based on a hidden Markov model to study various search configurations and to generalize the veto criteria by considering the overall veto performance in terms of efficiency and safety. The results can serve as a guideline for follow-up studies in future continuous gravitational-wave searches using a hidden Markov model algorithm. The results also apply qualitatively to other semicoherent search algorithms.

DOI: [10.1103/PhysRevD.106.123011](https://doi.org/10.1103/PhysRevD.106.123011)

I. INTRODUCTION

Gravitational waves (GWs), perturbations in spacetime which propagate at the speed of light, were first directly observed in 2015 when the Hanford and Livingston detectors of the Advanced Laser Interferometer Gravitational Wave Observatory (Advanced LIGO) detected a merging binary black hole system (GW150914) [1,2]. In the years since the first detection, the addition of Advanced Virgo [3] and the Kamioka Gravitational Wave Detector [4] to the network of observatories, together with improved sensitivity, has produced increasingly frequent detections of compact binary coalescences (CBCs) [5–7]. Other types of GW sources that are predicted to radiate at frequencies within the observational band of ground-based detectors remain undetected, e.g., the continuous gravitational waves (cw’s) produced by spinning neutron stars. Continuous gravitational waves, once

detected, will provide invaluable information regarding the structure of neutron stars as well as the nuclear equation of state [8]. A great deal of work has been carried out to develop methods to search for cw’s [9–11].

Since the expected strain amplitudes of cw’s are orders of magnitude smaller than those produced by CBCs, the computational cost of searching large template banks (including parameters such as the signal frequency and time derivatives thereof) coherently over a long period of time, e.g., ~ 1 yr, is high [9]. In addition, intrinsic, stochastic wandering of the frequency, sometimes called “timing noise” in the context of radio pulsars, could degrade the sensitivity of a cw search [12–14]. In this paper, we focus on a computationally efficient, semicoherent search strategy based on a hidden Markov model (HMM) which is equipped to track a signal frequency that wanders stochastically and spins down secularly [15,16]. The tracking scheme has its origins in engineering applications and has recently been used in many directed cw searches (e.g., [17–26]).

*Corresponding author.
dana.jones@anu.edu.au

The HMM-based searches are made up of two main procedures: (1) dividing the total observing time into subintervals and coherently calculating the signal power within each consecutive time interval (e.g., with length of ~ 1 d) using a frequency-domain matched filter (e.g., the \mathcal{F} -statistic) [27] and (2) using the Viterbi algorithm (a HMM tracking scheme) to find the most probable signal evolution over the total observing time (e.g., ~ 1 yr) [15,17]. This optimal signal evolution, referred to as the Viterbi path, consists of a frequency estimated at each discrete time step. The search over the full frequency band is usually parallelized and carried out in narrow subbands (with width of ~ 1 Hz). The Viterbi path obtained in each subband, corresponding to the most likely cw candidate in that subband, is assigned a Viterbi score. This score evaluates the candidate significance, such that a higher score signifies a greater probability that the candidate is inconsistent with random noise fluctuations [15,17]. For the mathematical details behind these procedures, see Secs. II A–II D in Ref. [22]. Some searches rely on an unnormalized log-likelihood instead of the Viterbi score to evaluate the significance of a candidate [20,23–26].

In this study, we focus on directed cw searches where the HMM-based methods have been most widely used. A directed search targets an astrophysical source at a known sky position with unknown ephemerides and hence is conducted over a wide frequency band (e.g., ~ 20 – 1000 Hz). A typical HMM-based directed search will produce on the order of 10^3 cw candidates, many of which will require the use of multiple verification techniques to be identified as noise. As such, a computationally efficient follow-up procedure is needed in order to comb through the results and distinguish the candidates caused by noise artifacts from any real astrophysical signal. This is perhaps even more true when considering an all-sky search, which is a survey done over the whole sky to look for cw signals and which typically produces on the order of at least 10^4 candidates that require follow-up (although some searches produce a list of candidates orders of magnitude larger, e.g., Einstein@Home searches [28]). As outlined in Refs. [29,30], a threshold is defined within these searches to obtain a reasonable total number of candidates in order to keep the follow-up computationally feasible. Defining more efficient candidate follow-up procedures would allow us to lower this threshold and process a greater number of candidates, thereby improving search sensitivity. This is clearly demonstrated in, e.g., Eq. (67) in Ref. [31], where the lowest strain amplitude that the search is sensitive to, denoted by $h_{0,\min}$, is directly related to the threshold CR_{thr} , such that a decrease in threshold would yield an improvement in sensitivity.

Many cw candidates, which have Viterbi scores above a threshold S_{th} (e.g., corresponding to a 1% false alarm probability), can be easily eliminated using already well-established procedures. Candidates are initially passed through a known-line veto and a single-interferometer veto, which are defined in Refs. [17,22–24]. The known-line veto involves eliminating every candidate whose frequency

evolution overlaps any known instrumental line present in any of the detectors [32,33]. The single-interferometer veto involves eliminating unidentified instrumental artifacts by checking if a candidate is significantly louder (i.e., if it has a larger Viterbi score) in one detector than in the other [34,35]. After this, most searches use two additional well-defined vetoes to further identify noise artifacts, also described in Refs. [17,22–24]. The total observing time can be split into multiple subintervals, and any candidates that are significantly louder in one particular subinterval are eliminated. Then, if the estimated frequency evolution rate is sufficiently low, the coherent time T_{coh} over which the data are integrated coherently can be increased and any candidate whose Viterbi score decreases when T_{coh} increases can be vetoed.

Although most candidates will have been eliminated after applying the aforementioned vetoes, some closely resemble astrophysical signals and require more careful inspection. Two additional veto strategies, both based on Doppler modulation (DM), have been developed for the candidate follow-up procedure [20,36,37]. They prove to be both useful and complementary to each other, in that one often vetoes the candidates that the other does not. In fact, the order of applying a series of vetoes is generally interchangeable. More computationally efficient vetoes are recommended to be used earlier. These two DM-based vetoes are more efficient compared to the vetoes that involve rerunning the search in subintervals or increasing the coherent time. Thus they are usually applied directly following the single-interferometer veto.

The DM-based vetoes are quite dependent on the search configuration used in a particular study. For the first of these two strategies, called the “DM-off veto,” the signal significance is calculated with the DM correction for the Doppler shift due to Earth’s motion being switched on and off. The candidates that are louder when the DM correction is switched off are not likely to be of astrophysical origin. For the second strategy, called the “off-target veto,” the search is shifted to one or more sky position(s) away from the true position of the target, i.e., a shifted DM correction is applied. If the candidates become louder, they are vetoed. Despite the wide applications of these two vetoes in existing searches [22–25], the general safety (i.e., no astrophysical signal is falsely eliminated) of these veto procedures still remains to be studied for various search configurations for HMM-based searches.

In this paper, we carry out an in-depth study of how the DM changes across the sky and investigate the impact on synthetic signals and noise outliers through simulations. In particular, we derive a set of rigorous criteria for when and how to use the off-target veto to discriminate between astrophysical signals and noise artifacts. We introduce a new concept of using an effective point spread function (EPSF) of the detection statistics obtained at various sky locations, which can serve as a particularly useful veto criterion. For the DM-off veto, which has been previously studied in coherent searches [36], we carry out additional

simulations to verify the criteria and safety in semicoherent searches. Although the simulations are designed for signals from isolated sources, the results can be applied to signals from cw sources in binary systems, assuming the Doppler shift caused by the orbital motion is fully accounted for. The impact from the imperfect removal of the binary orbital modulation is out of the scope of this paper. While this study is carried out using a HMM-based method, in principle the resulting guidelines broadly apply to stack-slide-based semicoherent algorithms [38]. Moreover, our results can be generalized to follow up candidates in an all-sky search.

The organization of this paper is as follows. In Sec. II, an analytic investigation of the Doppler modulation is presented. Section III outlines the search methods used in this study. Section IV introduces the EPSF and details the results of the off-target veto study. Section V describes the verification of the DM-off veto in semicoherent searches. A discussion of the results and concluding remarks are presented in Sec. VI.

II. DOPPLER SIGNATURE

In this section, we briefly review the cw signal model and calculate analytically how the DM affects cw signals. This provides a foundation for the empirical studies described in Secs. IV and V.

The phase of a cw signal can be modeled as follows [27]:

$$\Psi(t) = \Phi_0 + 2\pi \sum_{k=0}^s \frac{f_0^{(k)} t^{k+1}}{(k+1)!} + \frac{2\pi}{c} \mathbf{n}_0 \cdot \mathbf{r}_d(t) \sum_{k=0}^s \frac{f_0^{(k)} t^k}{k!}. \quad (1)$$

Here f_0 is the signal frequency at reference time $t = 0$ with respect to the solar system barycenter (SSB), $f_0^{(k)}$ is the k th time derivative of the instantaneous frequency evaluated at $t = 0$ at the SSB, \mathbf{n}_0 is the constant unit vector in the source

direction within the SSB reference frame, and $\mathbf{r}_d(t)$ is the position vector of the detector relative to the SSB origin. For the coordinate system with the SSB reference frame, we take the x axis parallel to the x axis of the celestial sphere coordinate system and the z axis perpendicular to the ecliptic z axis. Then, the unit vector \mathbf{n}_0 has the components [27]

$$\mathbf{n}_0 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \epsilon & \sin \epsilon \\ 0 & -\sin \epsilon & \cos \epsilon \end{pmatrix} \begin{pmatrix} \cos \alpha \cos \delta \\ \sin \alpha \sin \delta \\ \sin \delta \end{pmatrix}, \quad (2)$$

where ϵ is the angle between the ecliptic plane and Earth's equator (i.e., $\epsilon = 23.5^\circ$), α is the right ascension (RA) of the source, and δ is the declination (Dec). Meanwhile, the detector's position vector \mathbf{r}_d has the components [27]

$$\begin{aligned} \mathbf{r}_d = & R_{ES} \begin{pmatrix} \cos(\phi_0 + \Omega_0 t) \\ \sin(\phi_0 + \Omega_0 t) \\ 0 \end{pmatrix} \\ & + R_E \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \epsilon & \sin \epsilon \\ 0 & -\sin \epsilon & \cos \epsilon \end{pmatrix} \begin{pmatrix} \cos \lambda \cos(\phi_r + \Omega_r t) \\ \cos \lambda \sin(\phi_r + \Omega_r t) \\ \sin \lambda \end{pmatrix}, \end{aligned} \quad (3)$$

where $R_{ES} = 1$ AU is the mean distance between Earth's center and the SSB, R_E is the mean radius of Earth, Ω_0 is Earth's mean orbital angular velocity, Ω_r is Earth's rotational angular velocity, λ is the detector latitude, and ϕ_0 and ϕ_r are phases specifying the exact location of Earth in its orbital and diurnal motion, respectively, at $t = 0$.

Substituting Eqs. (2) and (3) into (1) yields the following expression [27]:

$$\begin{aligned} \Psi(t) = & \Phi_0 + 2\pi \sum_{k=0}^s \frac{f_0^{(k)} t^{k+1}}{(k+1)!} + \frac{2\pi}{c} \{ R_{ES} [\cos \alpha \cos \delta \cos(\phi_0 + \Omega_0 t) + (\cos \epsilon \sin \alpha \cos \delta + \sin \epsilon \sin \delta) \sin(\phi_0 + \Omega_0 t)] \\ & + R_E [\sin \lambda \sin \delta + \cos \lambda \cos \delta \cos(\alpha - \phi_r - \Omega_r t)] \} \sum_{k=0}^s \frac{f_0^{(k)} t^k}{k!}. \end{aligned} \quad (4)$$

We calculate the time derivative of $\Psi(t)$ to obtain the signal frequency

$$f(t) \approx (f_0 + f_0^{(1)} t) \left(1 + \frac{\mathbf{v}_d \cdot \mathbf{n}_0}{c} \right) \quad (5)$$

$$\begin{aligned} \approx & (f_0 + f_0^{(1)} t) \left\{ 1 + \frac{R_{ES} \Omega_0}{c} [-\cos \alpha \cos \delta \sin(\phi_0 + \Omega_0 t) \right. \\ & \left. + (\cos \epsilon \sin \alpha \cos \delta + \sin \epsilon \sin \delta) \cos(\phi_0 + \Omega_0 t)] \right\}, \end{aligned} \quad (6)$$

where $f_0^{(1)}$ is the first time derivative of the signal frequency at reference time $t = 0$ and \mathbf{v}_d is the velocity vector of the detector. In Eqs. (5) and (6), we omit the term corresponding to the DM effect due to Earth's rotation, which is negligible compared to the DM effect caused by Earth's orbital motion. We also omit the higher order time derivatives and other negligible terms in order to simplify calculations. (See Appendix A for the full derivation.) For convenience, we define the frequency shift induced by the DM as

$$\Delta \equiv \frac{\mathbf{v}_d \cdot \mathbf{n}_0}{c} \approx \frac{R_{ES}\Omega_0\kappa(\alpha, \delta, t)}{c} \quad (7)$$

with

$$\begin{aligned} \kappa = & -\cos\alpha \cos\delta \sin(\phi_0 + \Omega_0 t) \\ & + (\cos\epsilon \sin\alpha \cos\delta + \sin\epsilon \sin\delta) \cos(\phi_0 + \Omega_0 t). \end{aligned} \quad (8)$$

In Fig. 1(a), κ is plotted with $\phi_0 = 0$, $\Omega_0 = 2\pi/31557600 \text{ rad s}^{-1}$, and t set to Global Positioning System (GPS) time 1167545066, for a grid of evenly spaced RA and Dec values spanning the entire sky. Similar contour plots are computed for an equatorial coordinate

system in Ref. [39] and an ecliptic coordinate system in Ref. [40]. Figures 1(b) and 1(c) show two sets of curves where Δ is plotted as a function of time for an entire year (starting from GPS time 1167545066) and where each curve corresponds to a different sky position. These sky positions are marked on the contour map in Fig. 1(a) with white markers, each with a different shape that can be matched to its corresponding Δ curve (as a function of time) in Figs. 1(b) and 1(c).

We examine the DM of the signal frequency closely here in preparation for connecting the DM patterns plotted in Fig. 1 to the empirical results discussed in later sections of this paper. In particular, for any series of markers that lie

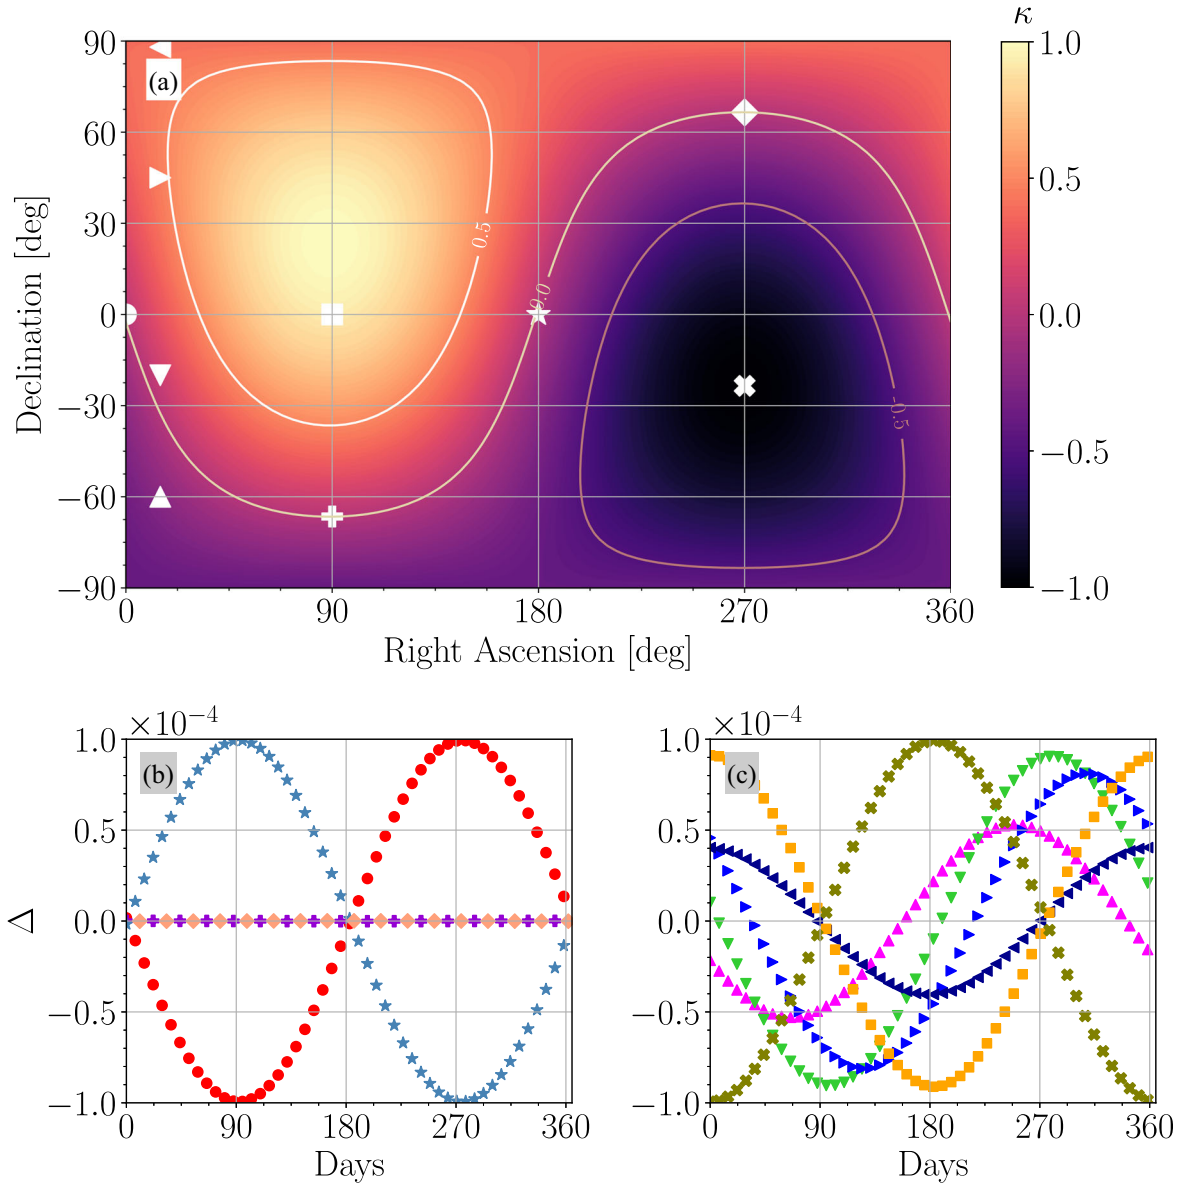


FIG. 1. Value of κ as a function of RA and Dec (both in degrees) for the GPS time 1167545066 (a). The Δ values as a function of time over a year (starting from 1167545066) for the sky positions highlighted in (a) are plotted in (b) and (c). Each curve, generated using one of these sky positions, is drawn with the same symbol as the white marker at this specific sky position in (a). All curves in (b) originate from the same contour in (a) with $\kappa = 0$. See more detailed discussions in Sec. IV; sky positions are given in Table III.

along the same κ contour in Fig. 1(a), the Δ curves that correspond to these markers will all intersect at the particular time used to calculate κ in (a). This is because the DM for sources at sky positions along a single contour are the same for this particular time. For instance, the Δ curves for four locations along the contour where $\kappa = 0$ are shown in Fig. 1(b), marked by red dots, light pink stars, green diamonds, and purple plus signs. The Δ value evolves over time in different ways depending on the sky position, so although the Δ curves all intersect at the start time, they no longer overlap as time shifts forward. While the dot and star markers are near the equator and experience strong DM effects, the diamond and plus sign markers are near the poles and are only weakly impacted by the DM. Examples of Δ curves for locations with a range of other κ values are displayed in Fig. 1(c).

For a source with a known sky position, Fig. 1 helps us to better understand how targeting a sky position shifted away from the true position of the source impacts the search results. That is, for all sky positions along a κ contour, despite being off target from the true position, the DM effect is the same as that seen at the true position for a given instant in time. As the integration time is increased, most of these similarities tend to disappear as the Doppler correction at each sky position evolves differently. However, the DM effects integrated over time for sky positions that lie along certain directions could still mimic each other, especially for sky positions that have similar κ values over the integration time. Therefore, when correcting the DM using an incorrect sky position, which is the basic principle behind the off-target veto, the recovered signal at the offset position could mimic the signal from the true sky position and produce similar detection statistics in a cw search. Indeed, if one searches a grid of locations around the true position of the source, the detection statistic for an elliptical area around the true sky position is higher than other sky regions. The nuances of this behavior are explained in Ref. [37] and discussed in detail, along with the different ways that the off-target veto has been implemented in recent searches to take advantage of this behavior, in Sec. IV.

III. GENERATING CANDIDATES

In order to study the cw validation techniques, we first generate candidates using the HMM-based semicoherent search method. The two fundamental procedures that make up the searches are as follows. First, we divide the observation time into subintervals and coherently calculate the signal power within each time segment using the maximum likelihood matched filter, \mathcal{F} -statistic. Second, we use a HMM tracking scheme to discover the most probable frequency evolution over the full observing time. These two procedures are briefly reviewed in Secs. III A and III B, respectively. The detection statistic used in this

study, known as the Viterbi score, is outlined in Sec. III C. More details can be found in Refs. [15,17].

A. \mathcal{F} -statistic

We can write the time-domain data collected in the detector as

$$x(t) = h(t) + n(t) \quad (9)$$

$$= \mathcal{A}^\mu h_\mu(t) + n(t), \quad (10)$$

where $h(t) = \mathcal{A}^\mu h_\mu(t)$ is the signal, and $n(t)$ is stationary, additive noise [27]. The amplitudes \mathcal{A}^μ , depending on the characteristic GW strain amplitude h_0 , source orientation, and initial phase, are associated with the four linearly independent components $h_\mu(t)$ that depend on the phase in Eq. (1) and the detector antenna patterns. (For more details, see Refs. [15,27]).

The \mathcal{F} -statistic is a frequency-domain matched filter that estimates the likelihood that a signal, parametrized by its frequency and the frequency time derivatives, is present in the data [27]. We first define a scalar product $(\cdot|\cdot)$ as a sum over single-detector inner products,

$$(x|y) = \sum_X (x^X|y^X) \quad (11)$$

$$= \sum_X 4\Re \int_0^\infty df \frac{\tilde{x}^X(f)\tilde{y}^{X*}(f)}{S_h^X(f)}. \quad (12)$$

Here X indexes the detector, x^X is the data in detector X , $S_h^X(f)$ is the single-sided power spectral density (PSD) of detector X , the tilde denotes a Fourier transform, and \Re is the real part of a complex number [41]. Using this definition, the \mathcal{F} -statistic can be expressed as

$$\mathcal{F} = \frac{1}{2} x_\mu \mathcal{M}^{\mu\nu} x_\nu, \quad (13)$$

where we define $x_\mu = (x|h_\mu)$, and $\mathcal{M}^{\mu\nu}$ is the matrix inverse of $\mathcal{M}_{\mu\nu} = (h_\mu|h_\nu)$ [42].

If we assume that the noise is Gaussian and that the single-sided PSD is the same in all detectors, the probability of having a signal in the data depends on the signal-to-noise ratio (SNR) when analyzing the data coherently for a time T_{coh} , given by (cf. Sec. III C in Ref. [27])

$$\rho_0^2 = \frac{K h_0^2 T_{\text{coh}}}{S_h(f)}, \quad (14)$$

where the constant K depends on the sky position, orientation of the source, and number of detectors.

B. HMM tracking and Viterbi algorithm

A general description of the HMM method can be found in Refs. [15,16]. We briefly summarize the method as follows.

A Markov chain is a stochastic process transitioning from one discrete state to another at discrete times $\{t_1, \dots, t_{N_T}\}$, where N_T is the total number of time steps. A HMM is made up of two variables, the unobservable, hidden state variable $q(t) \in \{q_1, \dots, q_{N_Q}\}$ and the observable, measurement state variable $o(t) \in \{o_1, \dots, o_{N_O}\}$, where N_Q and N_O are the total number of hidden and measurement states, respectively. The hidden state at time t_{n+1} is solely dependent on the state at time t_n (this is the Markovian assumption) and has a transition probability of

$$A_{q_j q_i} = P[q(t_{n+1}) = q_j | q(t_n) = q_i]. \quad (15)$$

At time t_n , the likelihood that the hidden state q_i is observed in state o_j is described by the emission probability, defined as

$$L_{o_j q_i} = P[o(t_n) = o_j | q(t_n) = q_i]. \quad (16)$$

The prior can be written as

$$\Pi_{q_i} = P[q(t_1) = q_i]. \quad (17)$$

Then, the probability that an observed sequence $O = o(t_1), \dots, o(t_{N_T})$ results from a hidden state path $Q = q(t_1), \dots, q(t_{N_T})$ via a Markov chain can be expressed as

$$P(Q|O) \propto L_{o(t_{N_T})q(t_{N_T})} A_{q(t_{N_T})q(t_{N_T-1})} \cdots L_{o(t_2)q(t_2)} \times A_{q(t_2)q(t_1)} \Pi_{q(t_1)}. \quad (18)$$

The most probable path, obtained by maximizing $P(Q|O)$, is simply [15]

$$Q^*(O) = \arg \max P(Q|O), \quad (19)$$

where $\arg \max(\dots)$ returns the argument that maximizes (\dots).

In this study, we track $q(t) = f(t)$, where $f(t)$ is the signal frequency at time t . We then map the discrete hidden states one-to-one with the frequency bins in the output of the \mathcal{F} -statistic calculated over a span of length T_{coh} (see Sec. III A), with each frequency bin size being $\Delta f = 1/(2T_{\text{coh}})$. Thus we choose T_{coh} to satisfy

$$\left| \int_t^{t+T_{\text{coh}}} dt' \dot{f}(t') \right| \leq \Delta f, \quad (20)$$

where $\dot{f}(t)$ is the first time derivative of the signal frequency. This relationship must remain valid throughout the total observing time T_{obs} .

The choice of $A_{q_j q_i}$ does not greatly impact the sensitivity of a HMM as long as it captures the general behavior of the signal, and as such, the particular transition matrix that is chosen should not impact the guidelines presented in this paper [16,43]. Without loss of generality, we assume that, in the cw signal we are searching for, the effect of timing noise on the frequency evolution is orders of magnitude smaller than that of the secular spin-down of the star, and that $|\dot{f}(t)|$ is uniformly distributed in the range from zero to the maximum estimated frequency derivative $|\dot{f}|_{\text{max}}$. It should be noted that, even if we are dealing with a situation where the timing noise is large and has a comparable effect on the signal's evolution as the star's spin-down, as long as an appropriate transition matrix is chosen based on our belief of the expected timescale on which the signal frequency evolves, the HMM remains capable of tracking the signal [15]. By substituting $|\dot{f}|_{\text{max}}$ into (20), we are able to simplify (15) to

$$A_{q_{i-1} q_i} = A_{q_i q_i} = \frac{1}{2}, \quad (21)$$

with all other $A_{q_j q_i}$ entries vanishing. (In searches where the signal frequency is assumed to walk randomly, such as in Ref. [16], we would instead have $A_{q_{i+1} q_i} = A_{q_{i-1} q_i} = A_{q_i q_i} = 1/3$.) The observed state $o(t)$ at time t is represented by the data observed by the detectors over $[t, t + T_{\text{coh}}]$. Using the definition of the \mathcal{F} -statistic, we express the emission probability as

$$L_{o(t)q_i} = P[o(t) | f_i \leq f(t') \leq f_i + \Delta f] \quad (22)$$

$$\propto \exp[\mathcal{F}(f_i)], \quad (23)$$

for $t \leq t' \leq t + T_{\text{coh}}$, where f_i denotes the central frequency in the i th bin. A uniform prior of $\Pi_{q_i} = N_Q^{-1}$ is chosen because there is no independent knowledge of the signal frequency at t_1 [15].

We use the classic Viterbi algorithm [44] to efficiently solve the HMM, which outputs the most likely frequency evolution path $Q^*(O)$ over the entire T_{obs} , i.e., the Viterbi path.

C. Detection statistic

We use the Viterbi score, denoted by S , to evaluate the significance of a candidate in this study [15,17]. In each subband searched (with width of 1 Hz in this paper), S is defined such that the log-likelihood of the optimal Viterbi path is equal to the mean log-likelihood of all paths ending in different bins of the subband plus S standard deviations at final step N_T . This is written as

$$S = \frac{\ln \delta_{q^*}(t_{N_T}) - \mu_{\ln \delta}(t_{N_T})}{\sigma_{\ln \delta}(t_{N_T})}, \quad (24)$$

with mean

$$\mu_{\ln \delta}(t_{N_T}) = N_Q^{-1} \sum_{i=1}^{N_Q} \ln \delta_{q_i}(t_{N_T}), \quad (25)$$

and variance

$$\sigma_{\ln \delta}(t_{N_T})^2 = N_Q^{-1} \sum_{i=1}^{N_Q} [\ln \delta_{q_i}(t_{N_T}) - \mu_{\ln \delta}(t_{N_T})]^2. \quad (26)$$

Here, $\delta_{q_i}(t_{N_T})$ is the likelihood of the path with the maximum probability ending in state q_i ($1 \leq i \leq N_Q$) at step N_T and $\delta_{q^*}(t_{N_T})$ is the likelihood of the optimal Viterbi path. In some applications of the HMM algorithm, e.g., [20,23–26], the total log-likelihood along the Viterbi path, $\ln \delta_{q^*}(t_{N_T})$, is directly used as the detection statistic. Transforming to the Viterbi score in Eq. (24) does not impact the results in this paper.

To conclude this section, it is worth noting that the \mathcal{F} -statistic is well equipped to deal with the DM correction; it calculates the expected Doppler shift for a given signal and thus corrects for the Doppler modulation when it computes the signal likelihood. The Viterbi score, which is calculated using the resulting likelihoods, is then also impacted by the DM correction. Then, the DM effect acts as an interesting basis for vetoes; it provides a useful way to study the difference between signals, whose intrinsic frequencies can only be recovered with the DM correction, and noise artifacts, whose frequencies are not impacted by the DM (since they originate on Earth).

IV. SEARCHING OFF TARGET

When an incorrect sky position is used to perform the DM correction on an astrophysical signal, the detection statistic decreases. The off-target veto makes use of this knowledge and of the fact that noise artifacts originate on Earth and are not impacted by DM due to Earth’s motion—that is, performing a DM correction on a noise artifact at a sky position that is slightly off target from the true sky position of the source should not cause the Viterbi score to drop significantly. In particular, we introduce a new concept of using an effective point spread function of the detection statistics obtained at various sky locations around the true position of the source to serve as a veto criterion in this section.

A. Implementation in other searches

In existing searches, the off-target veto is usually done as follows. The detection statistic is computed for one or more offsets from the source’s true sky position. For an astrophysical source, one would expect the detection statistic to drop and remain below a threshold once a certain offset from the true sky position is reached. On the other hand, the

detection statistic of candidates arising from noise artifacts should remain consistently above this threshold regardless of the offset [20]. Implementations of the off-target veto which are similar to the one considered in this study—namely, one or a handful of offsets is taken along one or two spatial directions—can be found in the following references: for a signal offset along one direction, see Refs. [20,25,45]; for two offsets, one along a line of constant RA and the other along a line of constant Dec, see Ref. [24]; and for multiple offsets along both a line of constant RA and a line of constant Dec, see Ref. [22].

However, this is not the only way the off-target veto is used in cw searches. In a recent search for cw’s from accreting millisecond x-ray pulsars, for each candidate that remains after the initial veto procedures, the log-likelihood is computed for a grid of off-target positions around the source’s true position [25]. Every candidate whose log-likelihood contours do not match simulations in Gaussian noise gets eliminated.

In Ref. [37], the authors use what they call the “sky-shifting” method where they run the search at many off-target sky positions across an entire hemisphere of the sky in order to build up an empirical background distribution of noise to be used as a detection statistic. Because an incorrect Doppler correction is being used for each of the off-target positions, the noise background is effectively blinded to the presence of any astrophysical signals. In addition to being more robust than simply choosing a handful of off-target sky positions, building up a background distribution is extremely useful in that it is better able to account for the non-Gaussianities found in real detector noise. However, the disadvantage here lies in the computational expense of generating the map for an entire hemisphere. Moreover, while this method would certainly be useful when targeting individual sources, it becomes computationally expensive when candidates are obtained for many sky positions, as in an all-sky search.

Another recent implementation of the off-target veto is presented in Ref. [40]. For an astrophysical signal, one can predict the relation between a mismatch in sky position and an error in estimated signal frequency and spin-down, using the DM relation indicated by Eq. (5). That is, for certain off-target sky positions, the incorrect Doppler correction still yields a significant detection statistic, but with some small quantifiable shift in the estimated frequency and frequency time derivative(s). Thus one can look for these expected relations, or “patterns,” that would be present were there a real signal. This method is robust in its ability to safely eliminate candidates caused by noise artifacts. However, it is more suitable for an all-sky search, where a group of candidates is usually found above the detection threshold across a patch in the sky and for a narrow band of estimated signal frequency and time derivative(s), than for a directed search. In addition, this method has only been

specifically applied in a Hough-transform-based pipeline (frequency Hough [31]).

Building upon these studies, robust and effective criteria are needed in HMM-based searches so that the off-target veto can be applied to a wide range of search configurations. In this section, we discuss an in-depth empirical study of the off-target veto and define veto criteria based on this investigation for two different scenarios depending on the number of candidates that require processing.

B. T_{coh}

Recalling Eq. (20), we can select a T_{coh} value that satisfies

$$T_{\text{coh}} \leq (2|\dot{f}_{\text{max}}|)^{-1/2}, \quad (27)$$

in a directed search, depending on a rough estimate of $|\dot{f}_{\text{max}}|$ for the source. Because T_{coh} directly depends on the rate of frequency evolution, it varies for searches targeting different sources. For example, for younger sources, a shorter T_{coh} is required [15]. In the first stage of the investigation, we test three different coherent lengths in order to help generalize the study to other stack-slide-based searches on various timescales as well as to allow for a broader application to other stack-slide-based semicoherent cw searches. We carry out the tests in both low- and high-SNR scenarios. As we increase T_{coh} , the recovered candidate is more significant at the position where the signal was injected, and the score drops below threshold more steeply as we move off target. This is exactly what we would expect to see; moreover, it holds true in both low- and high-SNR scenarios. The methods and results are detailed as follows.

The detection threshold S_{th} is estimated for each T_{coh} by running a series of 600 Monte Carlo simulations in pure Gaussian noise, with the detector's amplitude spectral density (ASD) set to $S_h^{1/2} = 4 \times 10^{-24} \text{ Hz}^{-1/2}$ here and throughout the rest of this study, in the subbands 200–201 and 500–501 Hz [46]. The resulting Viterbi scores are sorted and the score at the 99th percentile is set as the threshold (corresponding to a 1% false alarm probability per 1 Hz subband). For $T_{\text{coh}} = 12 \text{ h}$ and 5 d , we find $S_{\text{th}} = 5.56$ and 7.14 , respectively. (These values are cross-checked with previous studies using similar search configurations, e.g., Refs. [22,24]).

Next we run a series of simulations in which a synthetic signal is injected into Gaussian noise in the 1 Hz subband starting at 200 Hz and the Viterbi score is computed for a grid of offsets around the center (i.e., the sky position of the injection). Two signal strengths are studied: a weak signal with $h_0 = 4.0 \times 10^{-26}$ and a loud one with $h_0 = 8.0 \times 10^{-26}$. We assume the signals are circularly polarized with $\cos \iota$, where ι is the source inclination angle, fixed at unity such that h_0^{eff}/h_0 is held fixed, where h_0^{eff} is defined as [15]

$$h_0^{\text{eff}} = h_0 2^{-1/2} \{[(1 + \cos^2 \iota)/2]^2 + \cos^2 \iota\}^{1/2}. \quad (28)$$

When $\cos \iota$ is equal to 1 or -1 , the signal is circularly polarized and its strength is maximized. The simulation is run over 180 d (chosen to keep the computational cost of simulations to a reasonable level and given the fact that several recent cw searches use half-year observational data, e.g., [24,48–50]) starting from the arbitrarily chosen GPS time 1167545066 [same as the one used in Fig. 1(a)]. Three search configurations are tested in order to get a more complete picture of how a signal behaves around its center position in the sky: $T_{\text{coh}} = 12 \text{ h}$ ($N_T = 360$), $T_{\text{coh}} = 5 \text{ d}$ ($N_T = 36$), and $T_{\text{coh}} = 30 \text{ d}$ ($N_T = 6$). The fraction of the sky spanned by the grid of off-target positions varies depending on the signal SNR, where louder signals require a larger span in order for the extent of their features to be fully captured. However, to minimize the computational time needed for these simulations, a fixed grid of 21×15 data points is used regardless of the fraction of the sky spanned.

Figure 2 shows the Viterbi scores (plotted as the color) of the recovered signal at its injected sky position (RA = 00 h 00 m 00 s, Dec = 00° 00' 00") and for a grid of offsets around that position. The left and right columns correspond to the low- and high-SNR scenarios, respectively. The three rows, from top to bottom, correspond to $T_{\text{coh}} = 12 \text{ h}$, 5 d , and 30 d . Table I lists the Viterbi scores at the source's true position, denoted as S_{target} , for each panel.

The bright ellipse (with the brightest spot in the center) shown in these plots, which we refer to as the EPSF, is exactly what one would expect for an astrophysical signal in any stack-slide-based semicoherent search algorithm. (See Appendix B for an example using another stack-slide-based semicoherent search algorithm.) Moreover, such patterns are consistent with those in existing literature: e.g., see Fig. 1 in Ref. [51], Fig. 1 in Ref. [41], and Fig. 14 in Ref. [52]. The ellipse likely results from the different dependencies in RA versus Dec, outlined in Ref. [41], as well as the symmetries present in the Doppler modulation, as shown in Eq. (8) (discussed in Ref. [37]). The faint periodic features seen in some plots but not in others may also be related to the periodic functions within the Doppler modulation [40]. However, we rely only on the broader elliptical patterns in this study, so the minor features do not impact the results. A more detailed investigation of these minor features lies outside the scope of this paper. It should be noted that the ellipses shown in Figs. 2(a) and 2(b), both with $T_{\text{coh}} = 12 \text{ h}$, appear to be slightly off center such that the brightest point is not exactly at the injection site. This is most likely due to the relatively poor sky resolution when a short $T_{\text{coh}} < 1 \text{ d}$ is used. Indeed, the peak is recovered almost exactly at the injection location for $T_{\text{coh}} = 5 \text{ d}$ and 30 d .

In comparing the images shown in Fig. 2, it is clear that, as T_{coh} increases, the EPSF becomes brighter in the center




FIG. 2. Contour of Viterbi scores as a function of the offset in RA and Dec for an injection at RA = 00 h 00 m 00 s and Dec = 00° 00′ 00″ with signal strength $h_0 = 4.0 \times 10^{-26}$ (a),(c),(e) and $h_0 = 8.0 \times 10^{-26}$ (b),(d),(f). From top to bottom, coherent lengths $T_{\text{coh}} = 12$ h, 5 d, and 30 d are used with $T_{\text{obs}} = 180$ d. (Other simulation parameters: $S_h^{1/2} = 4 \times 10^{-24}$ Hz $^{-1/2}$, $\cos i = 1$.) See Table I for the Viterbi score obtained at the injection location (S_{target}) in each panel.

and narrower. In fact, it has already been established that the offset at which the detection statistic drops below threshold is related to T_{coh} ; as T_{coh} increases, this offset

TABLE I. Viterbi score obtained at the injection site for each panel shown in Fig. 2.

Figure 2	SNR	T_{coh}	S_{target}
(a)	Low	12 h	11.28
(b)	High	12 h	27.39
(c)	Low	5 d	62.70
(d)	High	5 d	176.59
(e)	Low	30 d	116.06
(f)	High	30 d	478.04

decreases [51]. This relationship is shown in Fig. 3; the offset in Dec (holding RA fixed at the injection coordinate) at which the Viterbi score drops below $0.5S_{\text{target}}$ is tracked as a function of T_{coh} . (To quantify this relation with better resolution, we conduct additional searches with two more coherent lengths $T_{\text{coh}} = 1$ d and 10 d in Dec.) The error bars for the offsets in Dec (too small to be shown in Fig. 3) are set by the grid spacing used to estimate these offsets, which is listed in Table II for each case, along with their errors ($\lesssim 1\%$). Indeed, for both SNRs tested, the shortest coherent time, 12 h, has the greatest Dec offset. As we increase T_{coh} , this offset decreases approximately linearly in log-log scale. We choose to take our offsets only in Dec because the behavior of the detection statistics is more

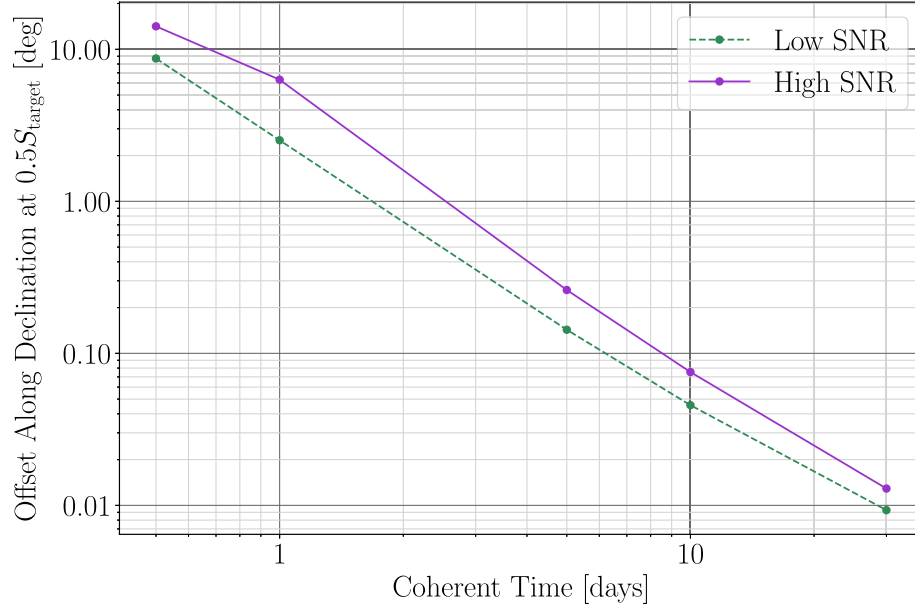


FIG. 3. Dec offset at which the Viterbi score drops below $0.5S_{\text{target}}$ as a function of T_{coh} . The coherent lengths tested are $T_{\text{coh}} = 12$ h, 1 d, 5 d, 10 d, and 30 d. The lines that join the sample points give a rough idea of where this offset would occur for other choices of T_{coh} within this range. For the low (green) and high (purple) SNRs, the signal strength of the injection is $h_0 = 4.0 \times 10^{-26}$ and $h_0 = 8.0 \times 10^{-26}$, respectively. (Other simulation parameters: $S_h^{1/2} = 4 \times 10^{-24}$ Hz $^{-1/2}$, $\cos i = 1$). See Table II for the offset grid spacing used for each configuration and the error on each point.

dependent on the source position when moving only in RA, so taking offsets in Dec allows us to set more generally applicable guidelines for the veto that are independent of sky position. (More details are discussed in Sec. IV C).

Figure 4 shows further details on how the Viterbi score decreases as the Dec offset increases for $T_{\text{coh}} = 12$ h, 5 d, and 30 d, from left to right. In general, in the first few steps away from the center, the scores drop steeply, especially for longer coherent times, then level out such that further increasing the offset leads to only small decreases in the score. Once a large enough offset is reached, the score drops below threshold and the signal is no longer detectable. This offset at which the signal becomes undetectable heavily depends on T_{coh} and SNR. For extremely loud

signals, the detection statistic may never drop below threshold, regardless of the sky position searched. However, current detectors do not operate in the high-SNR regime for cw sources—in fact, such a loud signal would be astrophysically implausible because the source would likely fall outside the range of extreme ellipticities and would be spinning down so rapidly that its \dot{f} value would not lie within the ranges used in this study. Thus we do not take into consideration the extremely high-SNR scenario.

C. Sky position

The elliptical EPSF shown in Fig. 2 is further explored in this section by varying the location of the injection. A series of simulations are run with different combinations of RA and Dec. We find that sky position is the dominant factor in determining the shape (i.e., the ratio of the major and minor axes) and orientation of the EPSF. The EPSF can act as a precise marker of an astrophysical signal, so we use it in the off-target veto criteria. The details are discussed below.

In this set of simulations, we fix the injection SNR using the signal strength $h_0 = 4.0 \times 10^{-26}$ and conduct the search with $T_{\text{coh}} = 5$ d. The other simulation parameters are the same as those in Sec. IV B. Figure 5 shows the resulting contour plots, where the Viterbi score (color) is plotted as a function of RA and Dec for a grid of offsets centered on the injection site; see Table III for the ten injection locations and their corresponding scores. Random variations aside, the shape and orientation of the EPSF change

TABLE II. Offset grid spacing used for each configuration and relative error on each point shown in Fig. 3.

SNR	T_{coh}	Grid spacing ($^\circ$)	Error (%)
Low	12 h	0.01	0.11
High	12 h	0.01	0.07
Low	1 d	0.01	0.40
High	1 d	0.01	0.16
Low	5 d	0.001	0.70
High	5 d	0.001	0.38
Low	10 d	0.0001	0.22
High	10 d	0.0001	0.13
Low	30 d	0.0001	1.08
High	30 d	0.0001	0.78




FIG. 4. Viterbi score (as a fraction of S_{target}) obtained at positions away from the center along Dec (from left to right: $T_{\text{coh}} = 12$ h, 5 d, 30 d). In each panel, orange dots and blue circles correspond to the low-SNR case with $h_0 = 4.0 \times 10^{-26}$ and high-SNR case with $h_0 = 8.0 \times 10^{-26}$, respectively. The solid orange and dashed blue lines mark the threshold, each one displayed as a fraction of S_{target} for its corresponding SNR. (Although S_{th} remains the same for a particular T_{coh} regardless of SNR, S_{target} obtained is larger for a higher SNR, and thus the orange and blue lines do not overlap).

systematically as functions of RA and Dec. Matching up each panel in Fig. 5 to its position in Fig. 1(a), we observe that the inclination of the ellipse at a particular sky position roughly follows the slope of the κ contour in Fig. 1(a) passing through that same position.

Two notable examples are discussed below. The first is presented in Fig. 5(h); this image is centered on RA = 01 h 00 m 00 s and Dec = $88^\circ 00' 00''$, and it shows an ellipse that spans more than two hours in RA—a significant fraction of the sky. Referring to where this position is found within the contour map in Fig. 1(a)—the triangle marker pointing toward the left—we immediately notice that the contour passing through this point moves along RA at a roughly fixed Dec. The same holds true for a source found at, for example, Dec = $-88^\circ 00' 00''$. Indeed, as the position approaches either pole, the EPSF extends broadly along RA because the DM effect is weak near the poles. This fact could lead to astrophysical signals being falsely vetoed as noise if the off-target veto is applied to sources near the poles using an offset only in RA. Certainly, choosing our criteria so that the off-target veto can be safely applied to candidates at the poles is important since it is unlikely that the DM-off veto will eliminate such candidates.

The second example worth discussing is shown in Fig. 5(j). This image is centered on RA = 18 h 00 m 00 s and Dec = $-23^\circ 30' 00''$ (the latter corresponding to the tilt of Earth’s axis) and shows an ellipse spanning roughly 15° in Dec. Once again, referring to the corresponding sky position in Fig. 1(a), we can see that the injection is located in the center of the dark region—the “dark spot” where $\kappa = -1$ (the cross marker). Based on the simulation results obtained at various positions and on how the EPSFs tend to vary with injection location, we find that this sky position at the dark spot—along with the “bright spot” with $\kappa = 1$ —in general produces the most extended EPSF in Dec. We use this to set veto guidelines in Sec. IV F.

Since the extension of the ellipse also largely depends on T_{coh} (for a given SNR), we further conduct a set of tests for injections at the dark spot, using another two choices of coherent length: $T_{\text{coh}} = 12$ h and 30 d. In Fig. 6, the top row shows the EPSFs at the dark spot for $T_{\text{coh}} = 12$ h (left), 5 d (middle), and 30 d (right). We inject a louder signal, $h_0 = 5.0 \times 10^{-26}$, for the search which uses $T_{\text{coh}} = 12$ h (as opposed to $h_0 = 4.0 \times 10^{-26}$ for the other two coherent times) so that the EPSF is better resolved. Of all the search configurations tested throughout this study, the ellipse at this sky position using the shortest coherent length ($T_{\text{coh}} = 12$ h) spans the largest angular distance in Dec—approximately 30° . The bottom row of Fig. 6 shows how the Viterbi score decreases as the Dec offset increases along the dashed cyan line plotted in the top panels, for $T_{\text{coh}} = 12$ h, 5 d, and 30 d, from left to right (similar to Fig. 4).

D. Real LIGO data in the second observing run

In addition to investigating the behavior of the EPSF of a signal in Gaussian noise, we look at some examples with a real noise background. The left panel in Fig. 7 shows the EPSF of a synthetic signal ($h_0 = 8.0 \times 10^{-26}$, $\cos i = 1$) injected into the 200–201 Hz band in the Advanced LIGO data collected in the second observing run (O2) [53], recovered using the HMM pipeline with $T_{\text{coh}} = 5$ d. The effective averaged ASD in that band is $S_h^{1/2} = 7 \times 10^{-24} \text{ Hz}^{-1/2}$. For consistency, a GPS start time of 1167545066 is used with a total observing time of 180 d. The right panel shows a signal with the same parameters injected into Gaussian noise ($S_h^{1/2} = 4 \times 10^{-24} \text{ Hz}^{-1/2}$). Despite the effects of non-Gaussianities in real detector noise and the higher noise ASD, the EPSF in LIGO O2 data is still clearly resolvable and, despite some additional features that seem to have caused a slight shift in the EPSF inclination, the overall

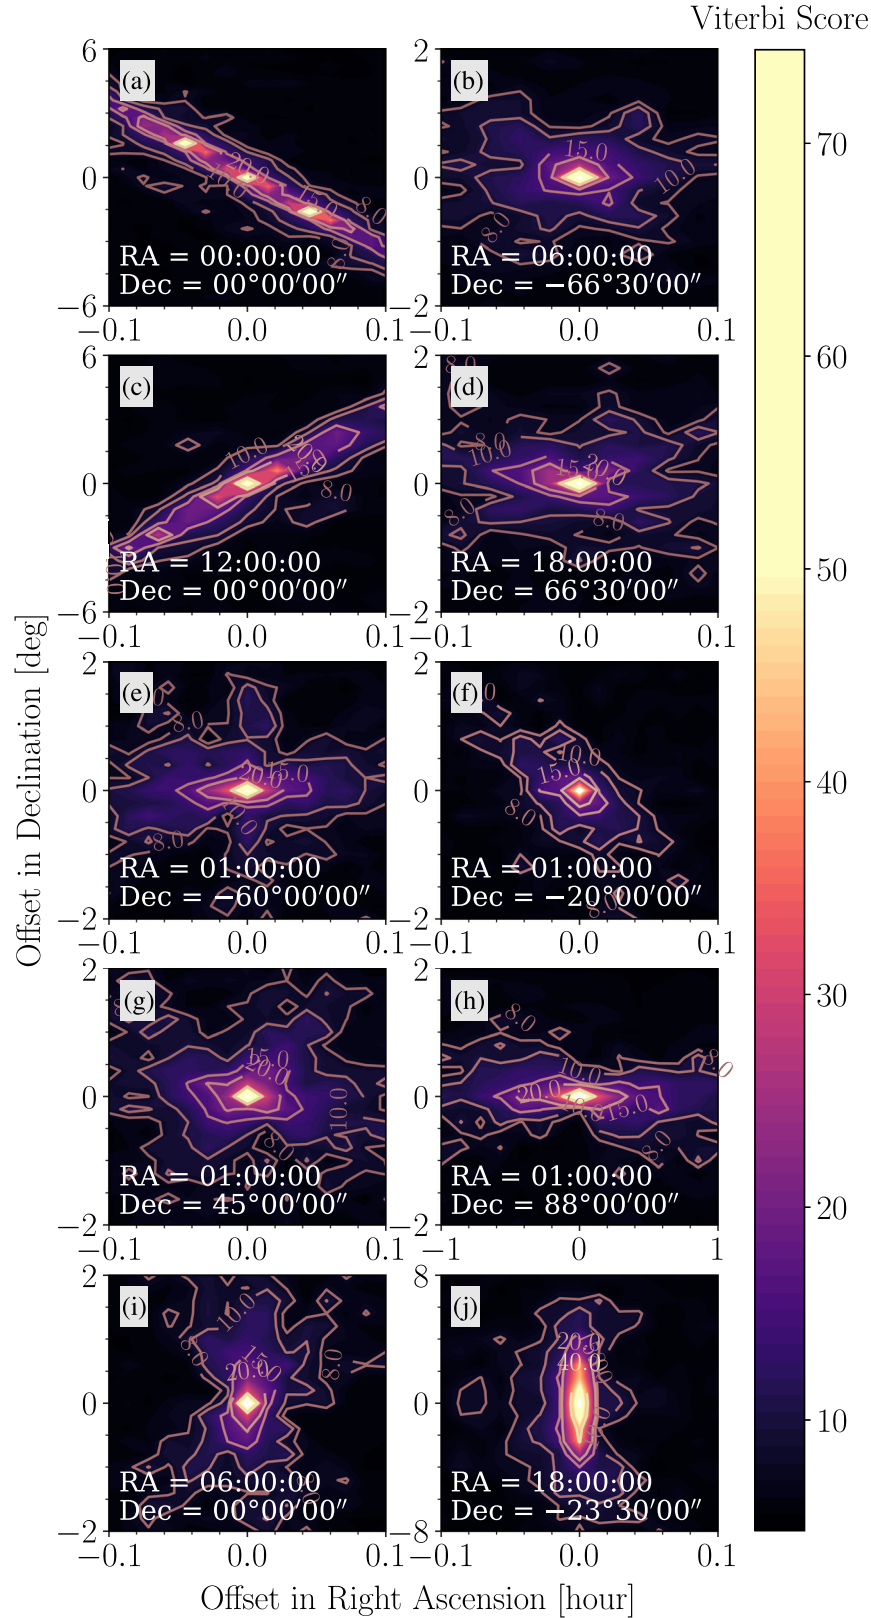


FIG. 5. Contour of Viterbi scores as a function of the offset in RA and Dec for injections at a variety of different sky positions with signal strength $h_0 = 4.0 \times 10^{-26}$ (Gaussian noise). Labels (a)–(j) correspond to the sky positions marked in the bottom left corners. A coherent length $T_{\text{coh}} = 5$ d is used with $T_{\text{obs}} = 180$ d. (Other simulation parameters: $S_h^{1/2} = 4 \times 10^{-24}$ Hz $^{-1/2}$, $\cos i = 1$.) See Table III for the Viterbi score obtained at each injection location (S_{target}).

TABLE III. Viterbi scores at the ten injection sites shown in Fig. 5.

Figure 5	RA	Dec	S_{target}
(a)	00:00:00	00°00'00''	62.70
(b)	06:00:00	−66°30'00''	73.20
(c)	12:00:00	00°00'00''	60.35
(d)	18:00:00	66°30'00''	73.71
(e)	01:00:00	−60°00'00''	70.93
(f)	01:00:00	−20°00'00''	55.82
(g)	01:00:00	45°00'00''	74.35
(h)	01:00:00	88°00'00''	69.42
(i)	06:00:00	00°00'00''	68.89
(j)	18:00:00	−23°30'00''	61.64

shapes of the EPSF are still comparable in these two sets of simulations.

Fifteen hardware injections were included in the LIGO O2 observing run for testing purposes [54]. These injections mimic real cw signals by moving the detector mirrors as though cw's are passing through, and they are an important complement to the software injections we have

looked at exclusively up until now. Of the 15 injections in O2, we look at two, P3 and P9, as examples. We have chosen these two injections because they have signal parameters that fall within the parameter space covered by a typical search configuration studied in this paper ($T_{\text{coh}} = 5$ d), enabling a consistent comparison and the use of statistics we have already derived here. We still search for $T_{\text{obs}} = 180$ d, starting from the GPS time 1167545066. The EPSFs of these injections are shown in Fig. 8, with P3 on the left and P9 on the right. P3 is located in the 108–109 Hz band and is centered at RA = 11 h 53 m 29.4 s and Dec = −33° 26' 11.8''. P9 is located in the 763–764 Hz band and is centered at RA = 13 h 15 m 32.5 s and Dec = 75° 41' 22.5''. See the Gravitational Wave Open Science Center O2 data release [53] for a more detailed parameter list. Overall, these EPSFs are in line with what we have predicted for astrophysical signals using software injections; if these injections were candidates in a real search, the off-target veto would not eliminate them (as demonstrated in Sec. IV F).

We also look at the Viterbi scores for a grid of sky positions centered at RA = 22 h 57 m 39.1 s and

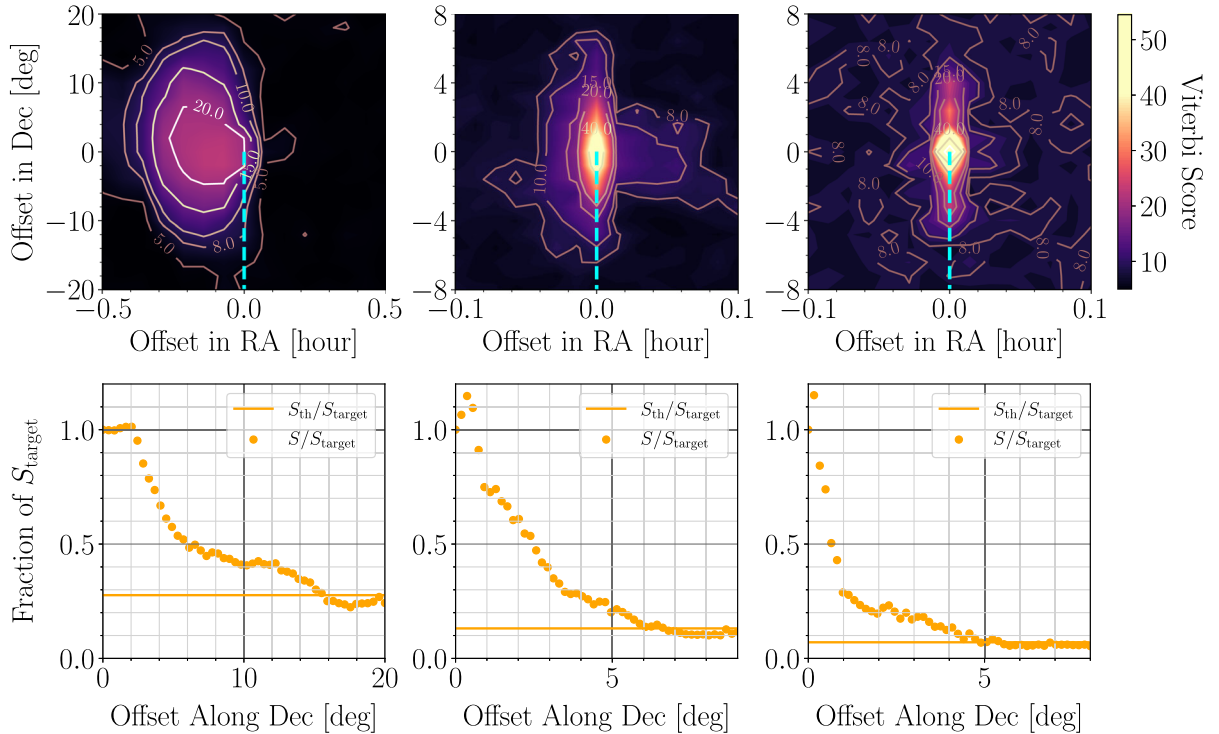


FIG. 6. Top row: contour of Viterbi scores as a function of the offset in RA and Dec for an injection at RA = 18 h 00 m 00 s and Dec = −23° 30' 00'' with signal strength $h_0 = 5.0 \times 10^{-26}$ (left) and $h_0 = 4.0 \times 10^{-26}$ (middle and right) (Gaussian noise), using three different choices of T_{coh} . From left to right, we have $T_{\text{coh}} = 12$ h, 5 d, and 30 d, with $S_{\text{target}} = 22.55$, 54.44, and 131.16, respectively. The total observing time is 180 d. (Other simulation parameters: $S_h^{1/2} = 4 \times 10^{-24}$ Hz $^{-1/2}$, $\cos i = 1$.) Bottom row: Viterbi score S (as a fraction of S_{target}) obtained at positions away from the center along Dec. In each panel, the search configuration is the same as that used in the contour image above it (from left to right: $T_{\text{coh}} = 12$ h, 5 d, 30 d). The solid lines mark the threshold for each search configuration displayed as a fraction of S_{target} . The dashed cyan line in each contour image in the top row marks the direction along which the series of offsets shown in the bottom row are taken.




FIG. 7. Contour of Viterbi scores as a function of the offset in RA and Dec for a synthetic signal ($h_0 = 8.0 \times 10^{-26}$, $\cos \iota = 1$) injected in LIGO O2 data (left, $S_h^{1/2} = 7 \times 10^{-24}$ $\text{Hz}^{-1/2}$) and Gaussian noise (right, $S_h^{1/2} = 4 \times 10^{-24}$ $\text{Hz}^{-1/2}$) in the 200–201 Hz band. The images are centered at the sky position of the injection (RA = 01 h 00 m 00 s and Dec = $-60^\circ 00' 00''$). A coherent length $T_{\text{coh}} = 5$ d is used with $T_{\text{obs}} = 180$ d.




FIG. 8. Contour of Viterbi scores as a function of the offset in RA and Dec for two hardware injections in LIGO O2 data. Left: the hardware injection P3 ($h_0 = 8.23 \times 10^{-25}$, $\cos \iota = -0.08$) in the 108–109 Hz band and centered at RA = 11 h 53 m 29.4 s and Dec = $-33^\circ 26' 11.8''$. Right: the hardware injection P9 ($h_0 = 3.01 \times 10^{-24}$, $\cos \iota = -0.62$) in the 763–764 Hz band and centered at RA = 13 h 15 m 32.5 s and Dec = $75^\circ 41' 22.5''$. We use $T_{\text{coh}} = 5$ d and $T_{\text{obs}} = 180$ d for both injections.

Dec = $-29^\circ 37' 20.0''$ in the 462–463 Hz subband in LIGO O2 data, where an unidentified noise artifact originating in the Hanford detector is located (identified in Sec. IV of Ref. [22]). The same duration of 180 d starting from 1167545066 is searched. Then we repeat this procedure in Gaussian noise with a synthetic signal ($h_0 = 2.0 \times 10^{-26}$, $\cos \iota = 1$) injected at RA = 22 h 57 m 39.1 s and Dec = $-29^\circ 37' 20.0''$. The EPSFs from LIGO O2 data and the simulated data containing a synthetic signal are shown in the left and right panels in Fig. 9, respectively. The left panel shows a spread of above-threshold Viterbi scores which fluctuate randomly across the sky. There is no trace of the bright peak in the center seen in the right panel resulting from a synthetic signal. This is the expected behavior of a candidate caused by a noise artifact, which would not be impacted by DM because it originates on Earth.

E. Physical insights from the EPSF

In practice, determining the exact shape and orientation of the EPSF of a cw signal is more complicated than simply following the direction of a Doppler pattern contour, e.g., in Fig. 1(a) (also see Ref. [40]). In fact, the shape and orientation of the EPSF mainly depend on a combination of three factors: the source position, the start time of the observation, and the amount of time the signal is integrated over. First, the sky position of the source is important, as this determines where it lies along the Doppler pattern shown in Fig. 1(a) and what Doppler correction should be made. Then, as the start time is shifted forward, this Doppler pattern also shifts across the sky. The EPSF becomes less extended as the observing time increases, because the bright and dark regions from different instants in time are included in the integration and




FIG. 9. Contour of Viterbi scores as a function of the offset in RA and Dec in the 462–463 Hz band in LIGO O2 data where an instrumental line lies (left) and Gaussian noise where a synthetic signal ($h_0 = 2.0 \times 10^{-26}$, $\cos i = 1$) is injected (right). Both images are centered at RA = 22 h 57 m 39.1 s and Dec = $-29^\circ 37' 20.0''$. A coherent length $T_{\text{coh}} = 5$ d is used with a total observing time $T_{\text{obs}} = 180$ d.

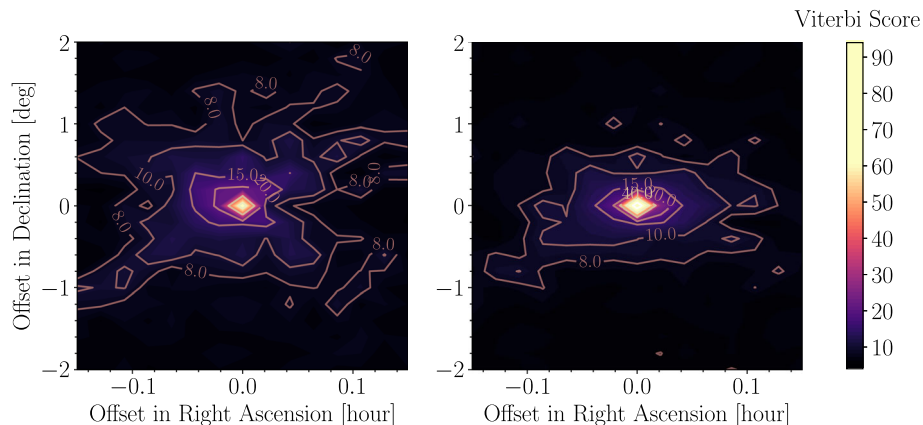


FIG. 10. Contour of Viterbi scores as a function of the offset in RA and Dec for an injection at RA = 18 h 00 m 00 s and Dec = $66^\circ 30' 00''$ with signal strength $h_0 = 4.0 \times 10^{-26}$ ($\cos i = 1$). Results are obtained for $T_{\text{obs}} = 180$ d (left) and $T_{\text{obs}} = 365$ d (right), with $T_{\text{coh}} = 5$ d for both. We find $S_{\text{target}} = 68.84$ (left) and $S_{\text{target}} = 93.94$ (right).

cancel out. This is demonstrated in Fig. 10, where the EPSF integrated over the full year (right) is brighter and more concentrated in the center than the EPSF integrated over half a year (left). Despite the impact from terrestrial noise, knowing these three factors allows one to roughly predict the shape and orientation of the EPSF around the true location in a real cw search, which can be used as a veto criterion.

F. Veto guidelines

To conclude this discussion of the off-target veto, we outline concrete guidelines for how to best apply this veto in two different scenarios: (1) when following up a large number of candidates from various sky positions and (2) when following up a handful of candidates from a limited number of positions. First,

in the follow-up procedure of either a directed search which is targeting many sources at different sky positions (e.g., [24,25]) or an all-sky cw search, if the earlier-stage vetoes are not effective enough, we need to use the off-target veto to process a large number of cw candidates from many sky positions. Thus, both safety and efficiency are important. Creating and inspecting detailed images of the EPSF for every candidate would be computationally expensive. This study allows us to use the detection statistic at a single Dec offset (as offsets in Dec have been shown in Sec. IV C to be safer than offsets in RA) to decide whether or not to veto a candidate. Moreover, we choose a conservative offset such that it is valid for all sky positions; however, the choice of offset does depend on the T_{coh} and T_{obs} used in a particular search.

1. Veto a large number of candidates

We propose the following broad criteria for applying the off-target veto to many candidates at once: a candidate can be safely vetoed *only if* the score at a certain conservative Dec offset remains above $0.5S_{\text{target}}$ or S_{th} (whichever is larger), with the added condition that the recovered Viterbi path at this offset position must overlap the original path. In practice, one can choose a different fraction other than 0.5 as needed, based on empirical studies. The reason we stipulate a fraction of the original score rather than a particular value of the score is so that, for a relatively loud signal with original score $S > 2S_{\text{th}}$, the dependence on the SNR is generally removed (as demonstrated in Fig. 3). If such an astrophysical signal exists, the score should drop at least 50% by the chosen Dec offset. We keep any candidate whose score is below S_{th} at this offset position because a weak signal becomes undetectable at very small offsets from the true sky position. One could even consider setting the offset based on a more sophisticated parameter-space-based distance, as presented in Ref. [55], but such a detailed study is beyond the scope of this paper.

Figure 6 shows an example of how one might find this optimal offset position for the three T_{coh} choices used most frequently in this paper. When processing many candidates at once, a plot like those along the bottom row of Fig. 6 can be created using a synthetic signal injected into Gaussian noise. To obtain a safe offset that can be applied to candidates from various sky positions, we carry out the simulations at the sky position where the detection statistic varies the slowest in Dec (i.e., the position at which the EPSF will be the most extended in Dec). The offset for applying the off-target veto can then be determined by looking at where the score drops below $0.5S_{\text{target}}$ for a relatively loud signal such that $0.5S_{\text{target}} > S_{\text{th}}$. For our particular search configuration, using the data presented in Fig. 6, we find this offset for $T_{\text{coh}} = 12$ h, 5 d, and 30 d (for $T_{\text{obs}} = 180$ d) to be 7° , 3° , and 1° , respectively. It should be noted for $T_{\text{coh}} = 12$ h that, despite the EPSF being significantly off center, as often occurs for short coherent times which have poorer resolution, the veto can still be safely applied as long as a conservative offset is chosen using the results from simulations in Gaussian noise.

The safety of this veto guideline is demonstrated by running a search for the two O2 hardware injections discussed in Sec. IV D (Fig. 8). Because both injections are very loud, we have $0.5S_{\text{target}} > S_{\text{th}}$. (In fact, the score at every location shown in both panels of Fig. 8 is above the detection threshold.) We veto the candidate if the score fails to drop below $0.5S_{\text{target}}$ at our chosen offset. Choosing to look at a single offset of 3° along Dec (as defined in the previous paragraph for this search configuration with $T_{\text{coh}} = 5$ d), we find that for both injections the score has dropped well below 50% of S_{target} . Thus, neither hardware injection is vetoed using the criteria we have defined.

2. Veto a handful of candidates

When we inspect a handful of candidates individually, a more comprehensive procedure is preferred, similar to what was done in Ref. [25]. That is, an image should be produced showing the Viterbi scores for a grid of sky positions centered at the candidate position. A second image also showing the EPSF, this time from a synthetic signal injected into Gaussian noise at the candidate sky position, should also be produced. The search configuration should remain unchanged when creating both images (i.e., observation start time, T_{coh} , and T_{obs} should all be the same), and the signal strength of the injection should be chosen such that the recovered signal at the center has roughly the same S_{target} as the cw candidate. Then, these two images can be compared. Although we would not expect the EPSF of a real cw signal to perfectly match the simulations, the elliptical pattern should, in general, be centered around the same position and have roughly the same shape and orientation. If it does not, this candidate may be safely vetoed.

Indeed, in Ref. [25], for the single candidate which remained after passing all the candidates through a hierarchy of vetoes, the candidate's EPSF in real data was not centered on the true sky position of the target and did not match the EPSF produced in Gaussian noise, so the candidate was vetoed. We verify this procedure in Fig. 9, where the Viterbi scores for a grid of sky positions centered at a candidate position in LIGO O2 data (vetoed as noise in Ref. [22]) and the EPSF for a synthetic signal injected into Gaussian noise at the same sky position are compared. No trace of the EPSF shown in the right panel can be seen in the real data in the left panel. Thus, as expected, we would veto this candidate.

It should be noted that, in a real search, data gaps are usually present (due to maintenance, upgrading, etc.). These gaps are generally dealt with by filling the emission probability in each frequency bin with equal probabilities for the periods during which no observational data are available, which will cause a slight decrease in overall sensitivity—akin to using a slightly shorter total observation time—but should not have any major impact on veto procedures such as the EPSFs. (See the sensitivity scaling due to gaps in the real interferometer data in, e.g., Fig. 2 of Ref. [17].) Nevertheless, for a more rigorous comparison, one can include data gaps in the Gaussian noise simulations.

V. SWITCHING OFF THE DM CORRECTION

Vetoes based on switching on and off the DM correction (i.e., the DM-off veto) have been studied in fully coherent searches [36] and used in existing studies [22–25]. In this section, we revisit the method and establish veto criteria in the context of semicoherent HMM-based searches. We verify that the veto is safe using Monte Carlo simulations.

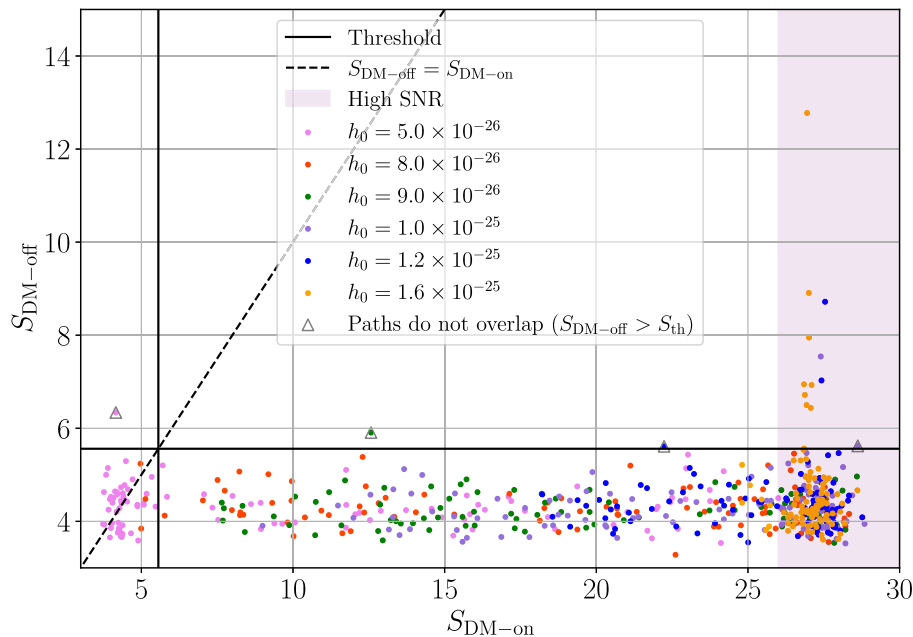


FIG. 11. Comparison of the Viterbi scores $S_{\text{DM-on}}$ and $S_{\text{DM-off}}$ for synthetic signals with $T_{\text{coh}} = 12$ h. Each color corresponds to a different signal strength, as indicated in the legend. The solid black lines mark S_{th} , the dashed black line marks the diagonal $S_{\text{DM-off}} = S_{\text{DM-on}}$, and the shaded purple region marks where the signal is so strong that it is detected in the search even without the DM correction applied. Any candidates marked with black triangles are not recovered in the DM-off search (the frequency paths in the DM-on and DM-off runs do not overlap) and thus can be ignored.

The DM-off veto, as first presented in Ref. [36], is a technique in which the DM correction, which accounts for the Doppler shift due to Earth’s orbital and rotational motion, is switched off and the detection statistic is reevaluated and compared to the one obtained with the DM correction applied. In a Viterbi search, the DM correction is switched off within the \mathcal{F} -statistic calculation. When the DM correction is switched off, the Viterbi score of a signal of astrophysical origin should drop below the threshold S_{th} and a different Viterbi path should be returned. A candidate resulting from a noise artifact, on the other hand, should yield a higher score (except in the rare case that a noise line wanders in a way that mimics the DM) and a Viterbi path that overlaps with the original. As such, a candidate is vetoed only if the Viterbi score increases when the DM correction is switched off and an overlapping Viterbi path is returned. In Ref. [36], the DM-off veto is described for a coherent \mathcal{F} -statistic search.

A. Search configurations

In order to safely use the DM-off veto in a semicoherent Viterbi search, we need to ensure that a single set of criteria can be applied to various search configurations.

Thus, we test two different T_{coh} values in our search, a longer one of 5 d and a shorter one of 12 h, and we use a total observation time of 180 d for the analysis (starting at the same arbitrarily chosen GPS time as in the off-target veto study).

B. Test synthetic signals

To establish criteria for the DM-off veto, we run pairs of simulations in which we inject synthetic signals into Gaussian noise and first search with the DM correction applied (DM on), then with DM off. We compare the Viterbi scores and paths between the pairs of simulations to see if the candidate present in the DM-on case disappears in the DM-off case, as would be expected for a real astrophysical signal. Figures 11 and 12 show the DM-off score ($S_{\text{DM-off}}$) plotted against the DM-on score ($S_{\text{DM-on}}$) for $T_{\text{coh}} = 12$ h and $T_{\text{coh}} = 5$ d, respectively, for randomly chosen sky positions distributed uniformly across the sky. A variety of different signal strengths are tested, ranging from $h_0 = 3.0 \times 10^{-26}$ to $h_0 = 2.4 \times 10^{-25}$, as shown in the legends. For each injected signal, we randomly draw $\cos i$ from a uniform distribution over the range $[-1, 1]$. In each figure, the horizontal and vertical black lines indicate the same detection threshold S_{th} for that particular T_{coh} .

The expected behavior of a synthetic signal is as follows: although $S_{\text{DM-on}} > S_{\text{th}}$ could lie anywhere along the horizontal axis to the right of the threshold depending on the SNR, $S_{\text{DM-off}}$ should fall below the dashed diagonal line (i.e., $S_{\text{DM-off}} < S_{\text{DM-on}}$). This behavior is confirmed in all of the injections. In fact, almost all of the markers fall below the horizontal black line marking S_{th} , with several exceptional cases. Among these exceptional candidates with $S_{\text{DM-on}} > S_{\text{DM-off}} > S_{\text{th}}$, a black triangle marks each




FIG. 12. Comparison of the Viterbi scores $S_{\text{DM-on}}$ and $S_{\text{DM-off}}$ for synthetic signals with $T_{\text{coh}} = 5$ d. Each color corresponds to a different signal strength, as indicated in the legend. The solid black lines mark S_{th} , the dashed black line marks the diagonal $S_{\text{DM-off}} = S_{\text{DM-on}}$, and the shaded purple region marks where the signal is so strong that it is detected in the search even without the DM correction applied. Any candidates marked with black triangles are not recovered in the DM-off search (the frequency paths in the DM-on and DM-off runs do not overlap) and thus can be ignored.

candidate whose Viterbi path recovered by the DM-off search does not overlap with the path returned by the DM-on search. When we describe the two paths as “overlapping” we mean that the DM-off path, widened by $\pm 1 \times 10^{-4} f_0$ Hz to account for the maximum Doppler shift, intersects the DM-on path at some point along its frequency evolution.

We now further discuss the two scenarios that result in the candidates produced in this study falling above the threshold in the DM-off search. The first is for those without triangle markers, all found at $S_{\text{DM-on}} \gtrsim 27$ for $T_{\text{coh}} = 12$ h ($h_0 \geq 1.0 \times 10^{-25}$) and $S_{\text{DM-on}} \gtrsim 200$ for $T_{\text{coh}} = 5$ d ($h_0 \geq 2.0 \times 10^{-25}$), shown within the shaded purple region which marks where the signals are so strong that we do not miss them even without the DM correction applied [24]. The signal is so loud in this case that it bleeds into nearby frequency bins within the \mathcal{F} -statistic, so the likelihood values for the bins around the true frequency bin are all high. When this occurs, even if we do not apply the DM correction in the \mathcal{F} -statistic, the total log-likelihood is still significant enough to produce a score above the threshold (but lower than the score with the DM correction switched on), and the Viterbi path still roughly tracks the true signal frequency. Nevertheless, we set the criteria to only veto candidates with $S_{\text{DM-off}}$ increased, which we do not see here, so the veto is still reliable for such loud signals.

The second scenario is for those weaker signals to the left of the shaded region but with $S_{\text{DM-off}} > S_{\text{th}}$. They are all marked by triangles, meaning that the DM-on and DM-off paths do not match each other. Given that we choose a threshold corresponding to a 1% false alarm probability per 1 Hz subband, these DM-off candidates are consistent with false alarms. We note that in both Figs. 11 and 12, one false alarm in the DM-off run (with $S_{\text{DM-off}} \approx 6.5$ and 7.0, respectively) does have an increased score relative to $S_{\text{DM-on}}$ that is above threshold (with a different Viterbi path), but in fact the injected signal is too weak to be identified as a candidate in the first place (i.e., $S_{\text{DM-on}} < S_{\text{th}}$). According to these simulation results, one can even safely veto a candidate if the score remains above S_{th} and the same Viterbi path is returned in the DM-off search, as long as the candidate is not in the high-SNR shaded region. [In practice, with the current detector sensitivity, candidates found in a real search with Viterbi scores $S > 27$ ($T_{\text{coh}} = 12$ h) and $S > 200$ ($T_{\text{coh}} = 5$ d) are caused by noise artifacts and would most likely be eliminated by another cw veto.] However, comparing $S_{\text{DM-off}}$ to $S_{\text{DM-on}}$ rather than the threshold is generally a more conservative option, i.e., to be cautious, we keep the candidates with $S_{\text{DM-on}} > S_{\text{DM-off}} > S_{\text{th}}$ for further scrutiny. In particular, setting veto criteria by comparing $S_{\text{DM-off}}$ to $S_{\text{DM-on}}$ rather than S_{th} ensures the veto safety for sources with sky positions close to

the ecliptic poles, since the DM correction at these positions would be minimal and so switching off the DM correction would not have much impact on the significance of the Viterbi score. Still, even for these sky positions, we would not expect the Viterbi score to increase when switching the DM correction off. Thus, based on the criteria we have defined, for the 1300 injections tested here, none are falsely eliminated and the DM-off veto remains safe.

C. Compare with noise

To provide a baseline for comparison, a series of synthetic monochromatic noise lines at fixed frequency (no DM is added in the simulation code) with different strain amplitudes h_{noise} are injected into the Gaussian noise background in the 200–201 Hz subband, and the results of the search are plotted in a similar fashion to the synthetic signals, shown in Fig. 13. Noise lines, unlike astrophysical signals, should increase in significance when the DM correction is switched off and lie above the diagonal (dashed line) in the figure. Indeed, all outliers caused by synthetic noise lines show this behavior other than the weakest few with $h_{\text{noise}} < 5.0 \times 10^{-26}$ that are below threshold for both the $S_{\text{DM-on}}$ and $S_{\text{DM-off}}$ searches (these exceptional ones are not identified as candidates in the first place).

The behavior caused by noise lines is also confirmed by studying the real noise lines from Advanced LIGO O2

(selected from the candidates identified as noise artifacts in Ref. [22]). The search results are shown in Fig. 14. Although quite a few candidates caused by noise lines lie below the dashed line, many do not have overlapping frequency paths between the DM-on and DM-off runs (signified with black triangles), meaning that the original candidates are not recovered by the DM-off runs, so they cannot be evaluated in the DM-off veto procedure. There are nine candidates caused by noise lines that fall below the dashed line that do have overlapping frequency paths, marked with red circles. These artifacts are inspected individually; Table IV presents a summary of their main attributes, namely the starting frequency for each artifact, which detector the artifact originates in, and at what stage in the veto hierarchy (as detailed in Ref. [22]) it is eliminated. There are a couple of theories as to why these noise artifacts are not eliminated by the DM-off veto: (i) the noise lines wander in a way that mimics the Doppler modulation in an astrophysical signal; or (ii) when we integrate data from both detectors and switch off the DM correction, the noise artifact at the candidate frequency remains the most significant, but the likelihoods in other frequency bins happen to increase, decreasing the relative significance of the original artifact. In a real search, these candidates would not be eliminated by the DM-off veto because they do not satisfy the veto criteria, so they would be followed up using other methods. In fact, as shown in Table IV, all nine candidates are eliminated by one of the two initial cw vetoes—the known-line veto, which identifies candidates




FIG. 13. Comparison of the Viterbi scores $S_{\text{DM-on}}$ and $S_{\text{DM-off}}$ for synthetic noise lines with $T_{\text{coh}} = 5$ d. Each color corresponds to a different injection strength, as indicated in the legend. The solid black lines mark S_{th} , and the dashed black line marks the diagonal $S_{\text{DM-off}} = S_{\text{DM-on}}$, above which we would expect the outliers caused by noise lines to lie.




FIG. 14. Comparison of the Viterbi scores $S_{\text{DM-on}}$ and $S_{\text{DM-off}}$ for real noise lines with $T_{\text{coh}} = 5$ d. The solid black lines mark S_{th} , and the dashed black line marks the diagonal $S_{\text{DM-off}} = S_{\text{DM-on}}$, above which we would expect the outliers caused by noise lines to lie. Candidates marked by black triangles are found at different frequencies in the DM-on and DM-off searches and thus cannot be regarded as a pair of results with and without the DM correction applied. Candidates marked with red circles have overlapping frequency paths with decreased $S_{\text{DM-off}}$ and are thus considered exceptional cases (see details in text).

as known instrumental lines, and the single-interferometer veto, which identifies candidates as previously unknown noise artifacts present in only one detector—so in all likelihood artifacts such as these would not even reach the stage wherein they would be analyzed using DM off. In summation, regardless of the small number of false positives (9 out of 108) when dealing with noise artifacts, the DM-off veto remains safe in that it does not falsely reject any of the 1300 synthetic signals we have tested above.

TABLE IV. Starting frequency and detector of origin for the nine noise artifacts in LIGO O2 data that the DM-off veto failed to veto (marked with red circles in Fig. 14). The last column lists the veto stage at which they are eliminated in the search presented in Ref. [22].

Starting frequency (Hz)	Detector	Veto stage
331.89	Hanford	Known line
444.55	Livingston	Known line
468.48	Livingston	Known line
471.50	Livingston	Known line
487.40	Livingston	Known line
492.06	Hanford	Single interferometer
504.80	Hanford	Known line
541.68	Hanford	Single interferometer
543.08	Hanford	Single interferometer

VI. CONCLUSION

In this study, we investigate the off-target veto and the DM-off veto in order to establish veto criteria for a semicoherent cw search. To aid the design of generalized and rigorous veto criteria for the off-target veto, we introduce an EPSF, discuss its physical insights, and demonstrate the applications. We draw conclusions about the safety of the DM-based vetoes through Monte Carlo simulations in which synthetic signals are injected into Gaussian noise. We use a combination of the coherent \mathcal{F} -statistic and a HMM scheme to track the signal, and we take the Viterbi score as our primary detection statistic.

We show that the off-target veto can be used in two different ways to follow up cw candidates depending on how many remain to be investigated. If many candidates from various sky positions need to be processed at once, a single Dec offset which is safe for all sky positions can be chosen such that, if the Viterbi score at this offset remains above 50% of the original score and above S_{th} , and the recovered Viterbi path overlaps the original candidate path, the candidate is vetoed. On the other hand, if only a few candidates require follow-up, a more detailed investigation can be carried out by calculating and comparing the EPSFs around the candidate location using both injections in Gaussian noise and the real data. A candidate is vetoed if these two patterns do not generally agree. Thus, the EPSF is used as a precise marker of an

astrophysical signal, expanding the applicability of the off-target veto.

Furthermore, we demonstrate that the DM-off veto is safe in the HMM-based semicoherent searches for all configurations tested (i.e., for all sky positions and for both a short coherent length of 12 h and a longer coherent length of 5 d). Although the veto is not able to eliminate every noise line, the majority are successfully eliminated. More importantly, no synthetic signals out of 1300 are falsely eliminated in any configuration tested.

With the conclusion of the third observing run of Advanced LIGO and Virgo and the impending fourth observing run, we will see continued improvements in detector sensitivity over the next few years. As a result, we expect to see searches produce more cw candidates (because the data are more likely to be polluted by weak noise artifacts), and these candidates will become more time consuming to verify or veto. This study verifies the safety of the off-target and DM-off vetoes and proposes a refined veto procedure in preparation for future cw searches and the first cw detection.

ACKNOWLEDGMENTS

We thank Evan Goetz for the helpful review and suggestions during the LIGO Publications and Presentations (P&P) review procedure. This material is based upon work supported by the U.S. National Science Foundation (NSF)'s LIGO Laboratory, which is a major facility fully funded by NSF. The authors are grateful for computational resources provided by the LIGO Laboratory and supported by NSF Grants No. PHY-0757058 and No. PHY-0823459. The authors acknowledge the support of the Australian Research Council Centre of Excellence for Gravitational Wave Discovery (OzGrav), Project No. CE170100004.

APPENDIX A: SIGNAL FREQUENCY

We simplify the cw phase shown in Eq. (4) by omitting the term that accounts for the DM effect due to Earth's rotation, which is 2 orders of magnitude smaller than the effect of Earth's orbital motion (i.e., $\Omega_0 R_{ES}/c \sim 1.0 \times 10^{-4}$, compared to $\Omega_R R_E/c \sim 1.5 \times 10^{-6}$). The approximate phase becomes

$$\begin{aligned} \Psi(t) \approx & \Phi_0 + 2\pi \sum_{k=0}^s \frac{f_0^{(k)} t^{k+1}}{(k+1)!} + \frac{2\pi}{c} \{R_{ES}[\cos \alpha \cos \delta \cos(\phi_0 + \Omega_0 t) + (\cos \epsilon \sin \alpha \cos \delta + \sin \epsilon \sin \delta) \sin(\phi_0 + \Omega_0 t)]\} \\ & \times \sum_{k=0}^s \frac{f_0^{(k)} t^k}{k!}. \end{aligned} \quad (\text{A1})$$

We can then write the signal frequency as

$$\begin{aligned} f(t) \approx & (f_0 + f_0^{(1)} t) + \frac{R_{ES}\Omega_0}{c} [-\cos \alpha \cos \delta \sin(\phi_0 + \Omega_0 t) + (\cos \epsilon \sin \alpha \cos \delta + \sin \epsilon \sin \delta) \cos(\phi_0 + \Omega_0 t)](f_0 + f_0^{(1)} t) \\ & + \frac{R_{ES}}{c} [\cos \alpha \cos \delta \cos(\phi_0 + \Omega_0 t) + (\cos \epsilon \sin \alpha \cos \delta + \sin \epsilon \sin \delta) \sin(\phi_0 + \Omega_0 t)] f_0^{(1)}, \end{aligned} \quad (\text{A2})$$

where we have omitted higher order frequency derivative terms $f_0^{(k)}$ with $k \geq 2$. Considering the fact that $\Omega_0 f_0 \gg f_0^{(1)}$ in the parameter space searched, the last term in Eq. (A2) can be omitted, and we obtain

$$f(t) \approx (f_0 + f_0^{(1)} t) \left\{ 1 + \frac{R_{ES}\Omega_0}{c} [-\cos \alpha \cos \delta \sin(\phi_0 + \Omega_0 t) + (\cos \epsilon \sin \alpha \cos \delta + \sin \epsilon \sin \delta) \cos(\phi_0 + \Omega_0 t)] \right\}, \quad (\text{A3})$$

i.e., Eq. (5) in Sec. II.

APPENDIX B: EPSF IN A SEMICOHERENT CROSS-CORRELATION SEARCH

To demonstrate that the EPSF and the off-target veto procedure can be generalized to other stack-slide-based semicoherent search methods, we run an analysis using

another method, the cross-correlation method (detailed in Refs. [56–58]). A similar EPSF image is produced to those output in the HMM-based analyses.

In a typical semicoherent search, detector data are divided into 30 min time segments and a short Fourier transform (SFT) is computed for each segment. When using the cross-correlation search pipeline, these SFTs are multiplied pairwise (where these pairs are chosen according to criteria

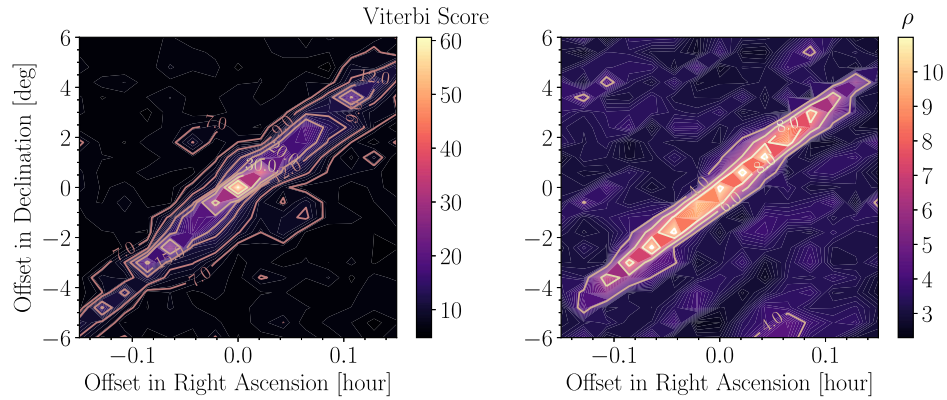


FIG. 15. Comparison of the EPSFs generated using the HMM method and the cross-correlation method. Left: contour of Viterbi scores as a function of the offset in RA and Dec for an injection in Gaussian noise at RA = 12 h 00 m 00 s and Dec = 00° 00' 00'' with signal strength $h_0 = 2.0 \times 10^{-26}$ ($\cos i = 1$) ($T_{\text{coh}} = 5$ d). Right: contour of ρ (used as the detection statistic in the cross-correlation method) as a function of the offset in RA and Dec for the same injection. (The maximum time lag is set to 1 h.) Both are for $T_{\text{obs}} = 180$ d.

that may involve incorporating data from multiple detectors and establishing a maximum time lag), yielding a cross-correlation variable for each SFT pair. The detection statistic ρ output by the search represents a weighted sum of these cross-correlation variables over all SFT pairs [58].

Figure 15 shows a side-by-side comparison of a synthetic signal with the same parameters (see figure caption)

injected into Gaussian noise and recovered using the HMM method (left) and the cross-correlation method (right). Although the two search methods use different detection statistics (S and ρ , respectively), they produce very similar EPSFs, demonstrating that the off-target veto methods detailed here are indeed generalizable to other semicoherent searches.

-
- [1] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Observation of Gravitational Waves from a Binary Black Hole Merger, *Phys. Rev. Lett.* **116**, 061102 (2016).
 - [2] J. Aasi *et al.* (LSC Collaboration), Advanced LIGO, *Classical Quantum Gravity* **32**, 074001 (2015).
 - [3] F. Acernese *et al.* (Virgo Collaboration), Advanced Virgo: A second-generation interferometric gravitational wave detector, *Classical Quantum Gravity* **32**, 024001 (2015).
 - [4] T. Akutsu *et al.*, Overview of KAGRA: Detector design and construction history, *Prog. Theor. Exp. Phys.* **2021**, 05A101 (2021).
 - [5] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo During the First and Second Observing Runs, *Phys. Rev. X* **9**, 031040 (2019).
 - [6] R. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run, *Phys. Rev. X* **11**, 021053 (2021).
 - [7] R. Abbott *et al.*, GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run, [arXiv:2111.03606](https://arxiv.org/abs/2111.03606).
 - [8] P.D. Lasky, Gravitational waves from neutron stars: A review, *Pub. Astron. Soc. Aust.* **32**, e034 (2015).
 - [9] K. Riles, Recent searches for continuous gravitational waves, *Mod. Phys. Lett. A* **32**, 1730035 (2017).
 - [10] 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).
 - [11] O. J. Piccinni, Status and perspectives of continuous gravitational wave searches, *Galaxies* **10**, 72 (2022).
 - [12] G. Hobbs, A. G. Lyne, and M. Kramer, An analysis of the timing irregularities for 366 pulsars, *Mon. Not. R. Astron. Soc.* **402**, 1027 (2010).
 - [13] 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).
 - [14] G. Ashton, D. I. Jones, and R. Prix, Effect of timing noise on targeted and narrow-band coherent searches for continuous gravitational waves from pulsars, *Phys. Rev. D* **91**, 062009 (2015).
 - [15] L. Sun, A. Melatos, S. Suvorova, W. Moran, and R. J. Evans, Hidden Markov model tracking of continuous gravitational waves from young supernova remnants, *Phys. Rev. D* **97**, 043013 (2018).
 - [16] S. Suvorova, L. Sun, A. Melatos, W. Moran, and Robin J. Evans, Hidden Markov model tracking of continuous

- gravitational waves from a neutron star with wandering spin, *Phys. Rev. D* **93**, 123009 (2016).
- [17] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Search for gravitational waves from Scorpius X-1 in the first Advanced LIGO observing run with a hidden Markov model, *Phys. Rev. D* **95**, 122003 (2017).
- [18] L. Sun and A. Melatos, Application of hidden Markov model tracking to the search for long-duration transient gravitational waves from the remnant of the binary neutron star merger GW170817, *Phys. Rev. D* **99**, 123003 (2019).
- [19] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Search for gravitational waves from Scorpius X-1 in the second Advanced LIGO observing run with an improved hidden Markov model, *Phys. Rev. D* **100**, 122002 (2019).
- [20] M. Millhouse, L. Strang, and A. Melatos, Search for gravitational waves from 12 young supernova remnants with a hidden Markov model in Advanced LIGO's second observing run, *Phys. Rev. D* **102**, 083025 (2020).
- [21] L. Sun, R. Brito, and M. Isi, Search for ultralight bosons in Cygnus X-1 with Advanced LIGO, *Phys. Rev. D* **101**, 063020 (2020).
- [22] D. Jones and L. Sun, Search for continuous gravitational waves from Fomalhaut *b* in the second Advanced LIGO observing run with a hidden Markov model, *Phys. Rev. D* **103**, 023020 (2021).
- [23] D. Beniwal, P. Clearwater, L. Dunn, A. Melatos, and D. Ottaway, Search for continuous gravitational waves from ten H.E.S.S. sources using a hidden Markov model, *Phys. Rev. D* **103**, 083009 (2021).
- [24] R. Abbott *et al.*, Searches for continuous gravitational waves from young supernova remnants in the early third observing run of Advanced LIGO and Virgo, *Astrophys. J.* **921**, 80 (2021).
- [25] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), Search for continuous gravitational waves from 20 accreting millisecond X-ray pulsars in O3 LIGO data, *Phys. Rev. D* **105**, 022002 (2022).
- [26] R. Abbott *et al.*, Search for gravitational waves from Scorpius X-1 with a hidden Markov model in O3 LIGO data, *Phys. Rev. D* **106**, 062002 (2022).
- [27] P. Jaranowski, A. Królak, and B. F. Schutz, Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection, *Phys. Rev. D* **58**, 063001 (1998).
- [28] B. Steltner, M. A. Papa, H.-B. Eggenstein, B. Allen, V. Dergachev, R. Prix, B. Machenschalk, S. Walsh, S. J. Zhu, O. Behnke, and S. Kwang, Einstein@Home all-sky search for continuous gravitational waves in LIGO O2 public data, *Astrophys. J.* **909**, 79 (2021).
- [29] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO and Advanced Virgo O3 data, [arXiv:2201.00697](https://arxiv.org/abs/2201.00697).
- [30] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), All-sky search for gravitational wave emission from scalar boson clouds around spinning black holes in LIGO O3 data, *Phys. Rev. D* **105**, 102001 (2022).
- [31] Pia Astone, Alberto Colla, Sabrina D'Antonio, Sergio Frasca, and Cristiano Palomba, Method for all-sky searches of continuous gravitational wave signals using the frequency-Hough transform, *Phys. Rev. D* **90**, 042002 (2014).
- [32] P. B. Covas *et al.* (LSC Instrument Authors), Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO, *Phys. Rev. D* **97**, 082002 (2018).
- [33] D. Davis *et al.*, LIGO detector characterization in the second and third observing runs, *Classical Quantum Gravity* **38**, 135014 (2021).
- [34] David Keitel, Reinhard Prix, Maria Alessandra Papa, Paola Leaci, and Maham Siddiqi, Search for continuous gravitational waves: Improving robustness versus instrumental artifacts, *Phys. Rev. D* **89**, 064023 (2014).
- [35] Paola Leaci, Methods to filter out spurious disturbances in continuous-wave searches from gravitational-wave detectors, *Phys. Scr.* **90**, 125001 (2015).
- [36] S. J. Zhu, M. A. Papa, and S. Walsh, New veto for continuous gravitational wave searches, *Phys. Rev. D* **96**, 124007 (2017).
- [37] Maximiliano Isi, Simone Mastrogiovanni, Matthew Pitkin, and Ornella Juliana Piccinni, Establishing the significance of continuous gravitational-wave detections from known pulsars, *Phys. Rev. D* **102**, 123027 (2020).
- [38] P. R. Brady and T. Creighton, Searching for periodic sources with LIGO. II. Hierarchical searches, *Phys. Rev. D* **61**, 082001 (2000).
- [39] R. Prix and Y. Itoh, Global parameter-space correlations of coherent searches for continuous gravitational waves, *Classical Quantum Gravity* **22**, S1003 (2005).
- [40] G. Intini, P. Leaci, P. Astone, S. D' Antonio, S. Frasca, I. La Rosa, A. Miller, C. Palomba, and O. Piccinni, A Doppler-modulation based veto to discard false continuous gravitational-wave candidates, *Classical Quantum Gravity* **37**, 225007 (2020).
- [41] R. Prix, Search for continuous gravitational waves: Metric of the multidetector F-statistic, *Phys. Rev. D* **75**, 023004 (2007).
- [42] C. Cutler and B. F. Schutz, Generalized F-statistic: Multiple detectors and multiple gravitational wave pulsars, *Phys. Rev. D* **72**, 063006 (2005).
- [43] B. G. Quinn and E. J. Hannan, *The Estimation and Tracking of Frequency* (Cambridge University Press, Cambridge, England, 2001), p. 266.
- [44] A. Viterbi, Error bounds for convolutional codes and an asymptotically optimum decoding algorithm, *IEEE Trans. Inf. Theory* **13**, 260 (1967).
- [45] H. Middleton, P. Clearwater, A. Melatos, and L. Dunn, Search for gravitational waves from five low mass x-ray binaries in the second Advanced LIGO observing run with an improved hidden Markov model, *Phys. Rev. D* **102**, 023006 (2020).
- [46] We use the software package LALSuite for all the simulations in this paper [47].
- [47] LIGO Scientific Collaboration, LIGO Algorithm Library—LALSuite, free software (GPL), [10.7935/GT1W-FZ16](https://doi.org/10.7935/GT1W-FZ16) (2018).

- [48] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), All-sky search for continuous gravitational waves from isolated neutron stars in the early O3 LIGO data, *Phys. Rev. D* **104**, 082004 (2021).
- [49] R. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), 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).
- [50] R. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Search of the early O3 LIGO data for continuous gravitational waves from the Cassiopeia A and Vela Jr. supernova remnants, *Phys. Rev. D* **105**, 082005 (2022).
- [51] S. Mastrogiovanni, P. Astone, S. D. Antonio, S. Frasca, G. Intini, I. La Rosa, P. Leaci, A. Miller, F. Muciaccia, C. Palomba, O. J. Piccinni, and A. Singhal, Phase decomposition of the template metric for continuous gravitational-wave searches, *Phys. Rev. D* **98**, 102003 (2018).
- [52] M. Isi, L. Sun, R. Brito, and A. Melatos, Directed searches for gravitational waves from ultralight bosons, *Phys. Rev. D* **99**, 084042 (2019).
- [53] Abbott R. *et al.*, Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo, *SoftwareX* **13**, 100658 (2021).
- [54] C. Biwer *et al.*, Validating gravitational-wave detections: The advanced LIGO hardware injection system, *Phys. Rev. D* **95**, 062002 (2017).
- [55] R. Tenorio, D. Keitel, and A. M. Sintes, Time-frequency track distance for comparing continuous gravitational wave signals, *Phys. Rev. D* **103**, 064053 (2021).
- [56] Sanjeev Dhurandhar, Badri Krishnan, Himan Mukhopadhyay, and John T. Whelan, Cross-correlation search for periodic gravitational waves, *Phys. Rev. D* **77**, 082001 (2008).
- [57] C. T. Y. Chung, A. Melatos, B. Krishnan, and J. T. Whelan, Designing a cross-correlation search for continuous-wave gravitational radiation from a neutron star in the supernova remnant SNR 1987A*, *Mon. Not. R. Astron. Soc.* **414**, 2650 (2011).
- [58] L. Sun, A. Melatos, P. D. Lasky, C. T. Y. Chung, and N. S. Darman, Cross-correlation search for continuous gravitational waves from a compact object in SNR 1987A in LIGO Science run 5, *Phys. Rev. D* **94**, 082004 (2016).