

Search for continuous gravitational waves directed at subthreshold radiometer candidates in O3 LIGO data

Alan M. Knee¹, Helen Du¹, Evan Goetz¹, Jess McIver¹, Julian B. Carlin^{2,3}, Ling Sun⁴, Liam Dunn^{2,3}, Lucy Strang², Hannah Middleton^{5,2,3} and Andrew Melatos^{2,3}

¹*Department of Physics and Astronomy, University of British Columbia, Vancouver, British Columbia V6T 1Z1, Canada*

²*School of Physics, University of Melbourne, Parkville, Victoria 3010, Australia*

³*OzGrav, University of Melbourne, Parkville, Victoria 3010, Australia*

⁴*OzGrav-ANU, Centre for Gravitational Astrophysics, College of Science, The Australian National University, Australian Capital Territory 2601, Australia*

⁵*Institute for Gravitational Wave Astronomy and School of Physics and Astronomy, University of Birmingham, Birmingham, B15 2TT, United Kingdom*



(Received 21 November 2023; accepted 23 February 2024; published 27 March 2024)

We present results of a follow-up search for continuous gravitational waves (CWs) associated with subthreshold candidates from the LIGO-Virgo-KAGRA (LVK) all-sky all-frequency (ASAF) directed radiometer analysis, using Advanced LIGO data from the third observing run (O3). Each ASAF candidate corresponds to a 1/32 Hz frequency band and ~ 13 deg² sky pixel. Assuming they represent possible CW sources, we analyze all 515 ASAF candidates using a semi-coherent, \mathcal{F} -statistic-based matched filter search. The search algorithm incorporates a hidden Markov model (HMM), expanding the signal model to allow frequency spin-wandering, as well as unmodeled frequency evolution of less than 10^{-5} Hz per day that is not captured by the searched range of $\pm 10^{-9}$ Hz/s in frequency derivative. Significance thresholds with a 5% probability of false alarm per ASAF candidate are determined empirically by searching detector noise at various off-target sky positions. We obtain 14 outliers surviving a set of vetoes designed to eliminate instrumental artifacts. Upon further investigation, these outliers are deemed unlikely to represent astrophysical signals. We estimate the sensitivity of our search to both isolated and binary sources with orbital period greater than one year by recovering simulated signals added to detector data. The minimum detectable strain amplitude at 95% confidence for isolated (long-period binary) sources is $h_0^{95\%} = 8.8 \times 10^{-26}$ (9.4×10^{-26}) at a frequency of 222.6 Hz. While this study focuses on ASAF subthreshold candidates, the method presented could be applied to follow up candidates from future all-sky CW searches, complementing currently existing methods.

DOI: [10.1103/PhysRevD.109.062008](https://doi.org/10.1103/PhysRevD.109.062008)

I. INTRODUCTION

Continuous gravitational waves (CWs) are an as-yet undiscovered class of persistent, quasimonochromatic gravitational waves (GWs) [1–4]. The canonical sources of CWs in the observing band of ground-based GW detectors are rapidly spinning, nonaxisymmetric neutron stars, which may be isolated or part of binary systems. Possible emission mechanisms include small deformations of the neutron star crust created by thermoelastic [5–8], magnetic [9–11], or tectonic [12–15] stresses, fluid oscillation modes (e.g., r -modes) that are unstable to gravitational radiation [16–18], and pulsar spin-up glitches [19,20]. A confirmed CW detection would offer novel insight into the physics of ultra-dense matter, allowing for new probes of the nuclear equation-of-state, measurements of neutron star ellipticities, and strong-field tests of gravity.

Current GW detectors, such as the Advanced Laser Interferometer Gravitational-wave Observatory (LIGO) [21], Advanced Virgo [22], and KAGRA [23], may be sensitive enough to detect CW sources within our galaxy. Searches for CW signals in detector data have targeted known young or millisecond pulsars [24–31], young supernova remnants [32–35], low-mass x-ray binaries [36–38], and the galactic center [39,40]. Various all-sky searches have also been carried out, scanning for CW emission from unknown neutron stars in isolated [41–45] or binary [46,47] configurations. Although there have been no CW detections reported thus far, improvements in detector sensitivity are allowing searches to probe deeper into the physically interesting regions of parameter space.

In this work, we carry out a search for CWs by following up subthreshold candidates identified by the LIGO-Virgo-KAGRA (LVK) all-sky all-frequency (ASAF) radiometer




FIG. 1. SNR, $\rho(f_0, \hat{n})$, versus central subband frequency, f_0 of the 515 subthreshold candidates from the LVK ASAF analysis, which are followed up in this work. The SNR is a function of f_0 and sky position, \hat{n} , where the latter symbol denotes a unit vector pointing at the candidate sky pixel. The bands around 500, 1000, and 1500 Hz are affected by optic suspension harmonics and have been notched out in the ASAF analysis.

analysis [48] (hereafter the “ASAF analysis”). While all-sky CW searches typically assume a deterministic signal coherent over $\sim 10^3$ s at minimum, and extending to even longer coherence times in the later follow-up stages [e.g., [44]], the ASAF analysis is an unmodeled search with comparatively weaker phase-coherence constraints. In the ASAF method, 192 s data chunks from pairs of detectors are frequency-binned and cross-correlated with sky-dependent phase and amplitude modulations, allowing one to search for a directional, narrowband stochastic signal. Adopting an equal-area tiling of the sky with 3072 pixels (≈ 13.43 deg² per pixel) and a coarse-grained 1/32 Hz frequency resolution, the ASAF analysis searched every frequency subband and sky-pixel pair from 20 Hz to 1726 Hz using LIGO-Virgo data from the first three observing runs (with the exception of a few bands excluded due to noise contamination). The ASAF analysis assumed a global signal-to-noise ratio (SNR) threshold with a 5% probability of false alarm over the full band. Although no candidates were found above this threshold, the ASAF analysis identified 515 subthreshold candidates (hereafter called “candidates”) with SNRs exceeding the 99th SNR percentile averaged over 10 Hz frequency bins [49]. Each candidate thus comprises a frequency subband and sky-pixel pairing, making it possible to conduct outlier follow-up with a CW search algorithm. The candidate SNRs and central subband frequencies are shown in Fig. 1. As the ASAF analysis can potentially detect coherent, quasimonochromatic signals, these candidates represent interesting targets for CW follow-up with more sensitive matched-filtering techniques. If any candidates are CW sources, then a matched-filter search with more demanding coherence requirements will improve the probability of detecting a signal. This is the first time CW analysis

techniques have been used to follow up candidates from a stochastic GW search.

Due to the model-agnostic approach of the ASAF analysis, the candidates could exhibit complicated or unpredictable frequency evolution, deviating from the canonical signal model assumed in many CW analyses. Several phenomena can give rise to such behavior, including stochastic spin-wandering, also known as *timing noise* [50–56], spin-up glitches [19,20], and long-period binary motion [57], all of which may degrade the SNR if left unaccounted for, though we leave a quantitative assessment of their impact on CW searches to future work. A signal of this kind could be missed in an all-sky CW search while being marginally detected in a stochastic search, where the assumption of a deterministic phase is relaxed. Motivated by the need for a more flexible signal model, we follow up the 515 ASAF candidates using a search algorithm which combines the semi-coherent \mathcal{F} -statistic [58,59] with a hidden Markov model (HMM), adapted from the method described in Refs. [60,61] and applied to various CW targets to date [e.g., [34–38,62–66]]. The HMM can track a signal as it wanders in frequency over the course of the observing period, broadening the CW signal model to include a wide variety of possible frequency evolution. We carry out this search on publicly available Advanced LIGO data from the third observing run (O3) [67–70].

This paper is organized as follows: The dataset used is described in Sec. II; in Sec. III, we review the implementation of our search algorithm; in Sec. IV, we describe our procedure for identifying CW outliers and screening likely instrumental artifacts; in Sec. V, we present the results of our follow-up searches and estimate the sensitivity; finally, we conclude in Sec. VI.

II. DATASET

We perform the search on O3 data [67] from the two Advanced LIGO detectors: LIGO Hanford (H1), and LIGO Livingston (L1). The O3 run began on April 1, 2019, 15:00 UTC (GPS time 1238166018.0), and ended on March 27, 2020, 17:00 UTC (GPS time 1269363618.0). The run was split into two parts, O3a and O3b, separated by a month-long commissioning period lasting from October 1 to November 1, 2019, when the detectors were nominally not in observing mode. During O3, the LIGO detectors operated at higher sensitivity and with a higher duty cycle than the previous O1 and O2 runs [71]. As including O1 and O2 data would increase the computational demands of the search without providing a significant enhancement of sensitivity, these data are not used. Similarly, we do not include data from Advanced Virgo or KAGRA due to their lower sensitivities [72,73].

The search algorithm processes a set of Tukey-windowed, short Fourier transforms (SFTs) from both detectors, taken over 1800 s of calibrated strain data [68–70,74]. We use the C01 calibration version O3 dataset [75,76], with loud noise

transients removed via a self-gating procedure [68] and 60 Hz power harmonics subtracted out [77]. Time segments determined to have serious data quality issues (Category 1 data quality vetoes) are not analyzed [69,78]. Any SFTs generated from segments with significant deadtime due to gating (30 s or longer) are also not used by our analysis. With these criteria, the total number of SFTs processed by our search is 11025 for H1 and 10133 for L1, covering 69.6% and 64.0% of the O3 run, respectively.

III. SEARCH METHOD

The search algorithm used to follow up each of the ASAF candidates consists of two main analysis steps. First, we filter the O3 data against a set of waveform templates covering the candidate subband and sky pixel, using the \mathcal{F} -statistic as the matched filter statistic. We then use the HMM formalism to reconstruct the optimal frequency path for each search template. We review the \mathcal{F} -statistic and HMM tracking technique in Secs. III A and III B, respectively. We construct template grids empirically for every candidate, as described in Sec. III C.

A. Semicoherent \mathcal{F} -statistic

The first step in our analysis is a matched filter search using a frequency domain estimator known as the \mathcal{F} -statistic. The \mathcal{F} -statistic is a maximum likelihood statistic for detecting CW emission from spinning neutron stars, modeled as triaxial ellipsoids which rotate about one of their principal axes [58,59]. It is denoted by $2\mathcal{F}(x|\lambda)$, where x represents the strain data, and $\lambda = \{\alpha, \delta, f, \dot{f}\}$ are the phase parameters of the filter template, consisting of a right ascension, α , declination, δ , GW frequency, f , and frequency derivative, $\dot{f} \equiv df/dt$. We do not include frequency derivatives beyond \dot{f} within the \mathcal{F} -statistic, but higher derivatives can be absorbed to some extent into the stochastic variations allowed by the HMM (see Sec. III B). Modulations of the signal frequency, caused by a nonzero \dot{f} as well as the motion of the Earth relative to the source, are accounted for within the \mathcal{F} -statistic calculation by demodulating the data with respect to the search template. Our pipeline leverages the `lalpulsar_ComputeFstatistic_v2` implementation of the multi-detector \mathcal{F} -statistic [70,79], which processes a set of SFTs as input and evaluates the matched filter on a grid of templates. The placement of templates for the ASAF candidates is described in Sec. III C.

We divide the O3 observing period into $N_T = 361$ contiguous intervals of duration $T_{\text{coh}} = 86400 \text{ s} = 1 \text{ d}$, and evaluate the \mathcal{F} -statistic coherently over each of these intervals. Phase information is thus preserved within each interval, but is otherwise not tracked between them, i.e., our search is semi-coherent. The choice of T_{coh} used in our search improves upon the 192 s coherence of the ASAF analysis, achieving a reasonable trade-off between

computational cost and sensitivity. Note that there are a total of 32 one-day segments from O3 with no SFTs from either detector, all but one being due to the O3 commissioning break. For days where no data are available, we substitute $2\mathcal{F} = 4$ uniformly for all frequencies [80], consistent with the expectation value of the \mathcal{F} -statistic in pure Gaussian noise. This procedure allows the HMM, described next in Sec. III B, to traverse data gaps in an unbiased manner, following previous CW HMM searches [e.g., [34,37]].

B. HMM tracking

The next step in our analysis is to find the optimal frequency path for each grid template $(\alpha, \delta, \dot{f})$ by means of HMM tracking [60,61]. HMMs are statistical models for inferring the behavior of an unobserved (“hidden”) state variable, $q(t)$, based on an observed state variable, $o(t)$, under the assumption that the observed state is influenced by the hidden state. Both states take on a sequence of values at N_T discrete epochs, $t_n \in \{t_1, \dots, t_{N_T}\}$. The hidden state is modeled as a stochastic Markov process, jumping between N_Q discrete values $q(t) \in \{q_1, \dots, q_{N_Q}\}$ with some transition probability matrix $A_{q_j q_i} = P(q(t_n) = q_j | q(t_{n-1}) = q_i)$. At each epoch t_n , the observed state $o(t_n)$ is related to $q(t_n)$ through the emission probability, $L_{o(t_n) q_i} = P(o(t_n) | q(t_n) = q_i)$.

Applied to CW analyses, $q(t)$ represents the heretofore unknown frequency of a CW signal, $f(t)$, which occupies one frequency bin q_i at any given epoch. The observed states, $o(t)$, represent the strain data collected in each of the N_T coherent intervals of a semicoherent search, where the interval duration is fixed by the coherence time, $T_{\text{coh}} = 1 \text{ d}$, of the matched filter. We use the \mathcal{F} -statistic, conditional on a $(\alpha, \delta, \dot{f})$ template, to map $o(t)$ onto $q(t)$, i.e., it is the emission probability for the HMM. The spin-wandering model is specified through the transition probability matrix, $A_{q_j q_i}$. As in previous HMM CW searches [34–38,60], we define the transition matrix as

$$A_{q_j q_i} = \frac{1}{3} (\delta_{q_j q_{i-1}} + \delta_{q_j q_i} + \delta_{q_j q_{i+1}}), \quad (1)$$

which restricts $q(t)$ to jump by 0 or ± 1 frequency bins $\Delta f = q_i - q_{i-1}$ over each epoch with equal probability. Since $q(t)$ can jump by no more than one frequency bin per coherent interval, the maximum amount of spin-wandering allowed by the HMM is limited to

$$|\dot{q}|_{\text{max}} = \Delta f / T_{\text{coh}}. \quad (2)$$

The probability that the signal follows some frequency path, $Q = \{q(t_1), \dots, q(t_{N_T})\}$, conditioned on the observed data, $O = \{o(t_1), \dots, o(t_{N_T})\}$, is given by




FIG. 2. Recovery of a simulated signal from a long-period binary system injected into three months of Advanced LIGO data from O3, using the search algorithm employed in this work to follow-up the ASAF candidates. The pixel colors indicate the value of the \mathcal{F} -statistic (i.e., emission probabilities) in each $T_{\text{coh}} = 1$ d interval (horizontal axis) and $\Delta f = 10^{-5}$ Hz frequency bin (vertical axis) after demodulating by the template \dot{f} and sky position. The frequency path recovered by the Viterbi algorithm is shown by the red curve. The binary orbit has eccentricity $e = 0.4$ with period $P_b = 10$ yr. The signal is circularly polarized with strain amplitude $h_0 = 8 \times 10^{-26}$.

$$P(Q|O) = \prod_{q(t_0)} \prod_{n=1}^{N_T} L_{o(t_n)q(t_n)} A_{q(t_n)q(t_{n-1})}, \quad (3)$$

where $\prod_{q(t_0)}$ is a prior probability on the initial frequency, $q(t_0)$, taken to be a uniform function of frequency within the search band, $\prod_{q(t_0)} \propto N_Q^{-1}$.

The objective of the HMM is to determine the optimal frequency path, $Q = Q^*$, that maximizes $P(Q|O)$. This is accomplished using the Viterbi algorithm—a dynamic programming algorithm which efficiently solves the HMM by strategically discarding suboptimal paths [81]. The detection statistic associated with Q^* (the Viterbi path) is the log-likelihood,

$$\mathcal{L} \equiv \log P(Q^*|O), \quad (4)$$

which is equivalent to incoherently summing \mathcal{F} -statistics along the frequency path, up to a linear offset. Although the Viterbi algorithm outputs N_Q optimal paths, each one terminating in a different frequency bin, we maximize over all such paths and report only the globally optimal path for each $(\alpha, \delta, \dot{f})$ template. In Fig. 2, we show an example of using the Viterbi algorithm, paired with the \mathcal{F} -statistic, to recover a simulated signal injected into O3-era detector noise. In this example, the source lies in long-period binary system, resulting in a slow Doppler modulation of the signal frequency which is successfully tracked by the Viterbi algorithm. We refer the reader interested in the full mathematical details of the Viterbi algorithm and HMMs to Ref. [60].

We have chosen to use $T_{\text{coh}} = 1$ d, which prevents us from tracking any frequency evolution happening on shorter scales. For reference, the low-mass x-ray binary Scorpius X-1 [82] is thought to exhibit relatively high amounts of spin-wandering, in which $T_{\text{coh}} = 10$ d has been

the typical spin-wandering timescale adopted by past Viterbi searches [36–38,80,83]. Thus, our Viterbi implementation has enough flexibility to capture the spin-wandering behavior of most known neutron stars.

C. Template grids

We search every ASAF candidate using a rectilinear grid of $(\alpha, \delta, \dot{f})$ templates spanning the candidate subband and sky pixel. This $1/32$ Hz subband is further subdivided into smaller frequency bins Δf . We then evaluate the \mathcal{F} -statistic for each grid template and Δf bin using the semi-coherent algorithm summarized previously in Sec. III A. We derive fixed grid resolutions empirically for each candidate, as described in Secs. III C 1 and III C 2 for frequency templates and sky templates, respectively. We opt to not use the standard parameter-space metric [59], as the spin-wandering allowed by the HMM signal model introduces additional correlations between \dot{f} and sky position templates. We verify that our template grid appropriately covers the required search space for each candidate in Sec. III C 3.

1. Frequency resolution

We divide each $1/32$ Hz candidate subband, centered on a frequency f_0 , into bins of width $\Delta f = 10^{-5}$ Hz. We search over various \dot{f} templates covering a $\pm 10^{-9}$ Hz/s range, with a spacing of $\Delta \dot{f} = 2 \times 10^{-10}$ Hz/s between templates. The bounds on the \dot{f} range correspond roughly to the maximum spin-down/up a signal could have and still be contained within one $1/32$ Hz ASAF subband. By searching over \dot{f} , the \mathcal{F} -statistic demodulation can correct for secular, long-term frequency trends over the observing period, leaving the HMM to track residual drifts occurring on timescales of order T_{coh} . Without including \dot{f} in the \mathcal{F} -statistic, the maximum frequency derivative recoverable by the HMM would be $|\dot{q}|_{\text{max}}$, requiring one to either decrease

T_{coh} or increase Δf to accommodate faster \dot{f} , either of which would make the search less sensitive. Combining the HMM with \dot{f} demodulation alleviates this restriction. We found that setting $\Delta \dot{f} \ll 2\Delta f/T_{\text{coh}}$ introduced correlations between adjacent \dot{f} templates, and hence chose $\Delta \dot{f} = 2 \times 10^{-10}$ Hz/s to mitigate these correlations. Finally, the \dot{f} values are defined with respect to a reference time, which we take to be the midpoint of O3 (GPS time 1253764818.0).

2. Sky resolution

The resolution of the sky grid depends primarily on the frequency and secondarily on the sky position of the ASAF candidate. We calibrate the sky grids by running the search algorithm on simulated signals and then examining the distribution of signal power around the injection point. We first estimate fiducial grid resolutions, $(\Delta\alpha_i^{\text{fid}}, \Delta\delta_i^{\text{fid}})$, across the entire sky for a fixed 1/32 Hz frequency band centered at 1 kHz, where i denotes the sky pixel index. We adopt the same HEALPix¹ [84,85] pixel basis as the ASAF analysis, so the pixel indices range from $i \in \{1, \dots, 3072\}$. After selecting a subset of 386 pixels uniformly spaced across the sky, we simulate a signal at the origin of each sample pixel with phase parameters $(f, \dot{f}) = (1 \text{ kHz}, 0 \text{ Hz/s})$, injecting into synthetic Gaussian noise SFTs created with `lalpulsar_Makefakedata_v5` [70] for H1 and L1. The generated SFTs include the same time gaps as the real O3 dataset, which is done to make the analysis as similar as possible to our O3 search. To ensure the signals can be confidently recovered, we fix $\sqrt{S_n}/h_0 = 1$, where h_0 is the signal amplitude and S_n is the one-sided power spectral density (PSD) of the detector noise. We search each pixel on a dense sky grid with $\Delta\alpha = \Delta\delta = 0.002$ rad using the search algorithm described in Secs. III A and III B. For these injections, we use the same frequency spacing given in Sec. III C 1 but do not search over \dot{f} , which is fixed at $\dot{f} = 0$. We compute the mismatch for each sky template as

$$\mu = 1 - \mathcal{L}/\mathcal{L}_0, \quad (5)$$

where \mathcal{L} is the Viterbi log-likelihood statistic [Eq. (4)] for the template and \mathcal{L}_0 is the log-likelihood of a template with identical parameters to the injection.

Figure 3 shows the mismatch contours surrounding an example injection, which lies at the origin of the diamond-shaped sky pixel. Contours corresponding to different mismatch levels are shown by the white curves. We calculate the fiducial sky resolutions $(\Delta\alpha_i^{\text{fid}}, \Delta\delta_i^{\text{fid}})$ for each sample pixel as follows:

$$\Delta\gamma_i^{\text{fid}} = \frac{1}{2} (\max_{\mu < 0.2} \gamma - \min_{\mu < 0.2} \gamma), \quad (6)$$

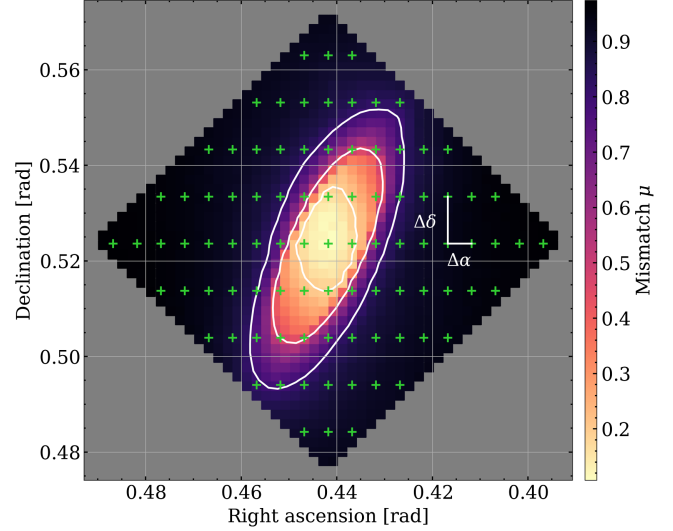


FIG. 3. Example of a fiducial sky grid constructed using a simulated signal at $f_0 = 1$ kHz, $\alpha_0 = 0.442$ rad, $\delta_0 = 0.524$ rad. The diamond-shaped search region is the sky pixel. The mismatch as a function of sky position, μ , is indicated by the color scale. The $\mu = 0.2, 0.5$, and 0.8 contours are shown by the white outlines. The green crosses show the resulting sky grid estimated from the $\mu = 0.2$ contour.

where $\gamma \in \{\alpha, \delta\}$, i.e., we estimate the α and δ ranges of the sky region enclosed by the $\mu = 0.2$ contour, and take half those ranges as the sky resolution. This yields a coarser search grid, as shown by the green crosses in Fig. 3, which is built starting from the pixel origin such that a template always lands on the origin. The $(\Delta\alpha_i^{\text{fid}}, \Delta\delta_i^{\text{fid}})$ values are then interpolated over all 3072 possible sky pixels. The sky grid for a particular ASAF candidate located at pixel i and central frequency f_0 is found through a simple scaling of this fiducial grid,

$$\Delta\gamma_i = \Delta\gamma_i^{\text{fid}} \left(\frac{1 \text{ kHz}}{f_0} \right). \quad (7)$$

The computational cost of the search scales with the square of the frequency, owing to the greater number of sky templates needed to guarantee a reasonable mismatch. The number of searched sky templates, N_{sky} , for all ASAF candidates is shown in Fig. 4, and ranges from a single sky position for $f_0 < 120$ Hz, up to $N_{\text{sky}} \sim 10^3$ for the highest-frequency candidates.

3. Mismatches

We verify that our template grids provide an acceptable range of mismatches by simulating the search on 200 mock ASAF candidates with random sky pixels and 1/32 Hz frequency subbands. For each mock candidate, we perform 20 separate injection-recovery tests, in which a signal is injected into synthetic Gaussian noise in H1 and L1 (fixing

¹<http://healpix.sourceforge.net>.




FIG. 4. Number of searched sky templates, N_{sky} , versus central subband frequency for all ASAF candidates. The number of templates scales with frequency squared, but also varies with the sky position. The total number of $(\alpha, \delta, \dot{f})$ templates is given by $N_{\text{tot}} = N_{\text{sky}}N_{\dot{f}}$, where $N_{\dot{f}} = 11$ for all candidates.

$h_0/\sqrt{S_n} = 1$) with frequency and sky position sampled uniformly within the subband and sky pixel, and \dot{f} sampled uniformly from $\pm 10^{-9}$ Hz/s. We analyze the simulated data twice—first by searching with the empirical template grid, and then searching at a template with identical parameters to the injection, which lets us compute a mismatch via Eq. (5). To reduce the computational burden, we only search templates within one parameter grid step from the injection.² The probability distribution function of the mismatch, $p(\mu)$, between the most significant grid template and the injection-matching template is shown in Fig. 5. The median is $\mu = 0.135$, and is constrained below $\mu = 0.386$ at the 95th percentile. Although the sky grids are constructed via the $\mu = 0.2$ mismatch contour, the empirical mismatch distribution extends to $\mu > 0.2$ at the tail, which is due to the additional parameter space complexity induced by searching over both \dot{f} and sky position.

IV. OUTLIER IDENTIFICATION

In this section, we outline our procedure for identifying interesting outliers resulting from our ASAF follow-up. We set empirically derived thresholds for every ASAF candidate by searching detector data at numerous sky positions shifted off-target from the ASAF sky pixel, following the procedures in Refs. [37,38]. Shifting the search off-target allows one to sample the distribution of the detection statistic, $p(\mathcal{L})$, in noise, i.e., not in the presence of a

²Rarely, when the candidate is at low frequency ($f_0 \lesssim 100$ Hz) and located near the celestial poles, and a signal is injected near the edge of the sky pixel (inclusive), there are no templates within one grid step of the injection. This occurs because of the peculiar pixel shapes near the poles and down-scaled template density from searching at low frequency. In such cases, we expand the search to include templates within two grid steps from the injection.

putative signal. The thresholds, \mathcal{L}_{th} , are calculated based on a pre-determined false alarm probability, which we set at 5% per ASAF candidate, i.e., each candidate has a 5% chance of yielding on outlier due to random noise. For each candidate, we record the loudest log-likelihood template, $\mathcal{L}^* \equiv \max \mathcal{L}$, and compare it to the threshold. If $\mathcal{L}^* > \mathcal{L}_{\text{th}}$, then the template is deemed an “outlier” and is flagged for further investigation.

A. Thresholds

Historically, HMM CW searches have set thresholds using one of two methods: a parametric approach, where an exponential model is fitted to the tail of the noise distribution [80,86], and a nonparametric approach, where the threshold is defined to be some percentile of the distribution [e.g., [38]]. While the latter method is more robust to template correlations and non-Gaussian data, it generally requires a much larger number of off-target simulations to guarantee an accurate percentile. Due to the large computational cost of the number of trials required for the percentile method, we employ the parametric approach, which reduces the computational burden.

Following Ref. [37], we model the tail of the off-target distribution as

$$p(\mathcal{L}) = k\lambda e^{-\lambda(\mathcal{L}-\mathcal{L}_{\text{tail}})}, \quad \mathcal{L} > \mathcal{L}_{\text{tail}}, \quad (8)$$

where λ and $\mathcal{L}_{\text{tail}}$ are the slope and location parameters, respectively, and k is a normalization factor. Because the detector sensitivity varies as a function of frequency and sky position, we expect the thresholds to vary depending on the location of the ASAF candidate. Thus, we sample and fit [using Eq. (8)] the off-target distributions for all ASAF candidates, deriving for each candidate an independent threshold. In order to obtain samples from the off-target distribution, we search at $N_{\text{off}} = 4000$ off-target positions from a disjoint sky region around each ASAF candidate,

$$\alpha \in [\alpha_{\text{min}} - 20^\circ, \alpha_{\text{min}} - 5^\circ] \cup [\alpha_{\text{max}} + 5^\circ, \alpha_{\text{max}} + 20^\circ];$$

$$\delta \in [\delta_{\text{min}}, \delta_{\text{max}}], \quad (9)$$

where α_{min} , α_{max} (δ_{min} , δ_{max}) are the minimum and maximum right ascension (declination) values of the template grid. A buffer of 5° in right ascension is included to mitigate contamination of the off-target samples by a possible signal in the on-target sky pixel. Every off-target point is searched with the same frequency binning and \dot{f} grid as the on-target search. With $N_{\text{off}} = 4000$ off-target points, we obtain $M = N_{\text{off}}N_{\dot{f}} = 44000$ sample statistics, \mathcal{L}_i , per candidate. The cutoff for the start of the tail, $\mathcal{L}_{\text{tail}}$, is defined as the 96th percentile of the off-target distribution, by which point the tails consistently appear exponential as long as the candidate subband does not contain any loud




FIG. 5. Probability distribution function (PDF), $p(\mu)$, and cumulative distribution function (CDF) of the mismatch, obtained using 4000 injection-recovery tests to validate the template grids. The PDF is represented by the gray histogram bins with the scale given on the left axis, while the CDF is the solid black curve with the scale on the right axis. The solid and dashed vertical lines indicate the median and 95th percentile mismatch values, respectively. The tail of larger mismatches is due to circumstances in which the injected signal simultaneously falls far from both a \hat{f} and sky position template.

non-Gaussian disturbances.³ This leaves $N_{\text{tail}} = 0.04M = 1760$ samples in the tail to estimate λ ($k = 0.04$ is already fixed by the tail cutoff), for which we use the maximum likelihood estimator

$$\hat{\lambda} = \frac{N_{\text{tail}}}{\sum_{i=1}^{N_{\text{tail}}} (\mathcal{L}_i - \mathcal{L}_{\text{tail}})}. \quad (10)$$

An example fit to the off-target distribution using Eq. (10) is shown in Fig. 6. In the Gaussian limit, the expected 2σ relative error in $\hat{\lambda}$ is $2N_{\text{tail}}^{-1/2} \sim 4.7\%$, which is also shown in Fig. 6 by the shaded region around the best-fit line. This is in agreement with Monte Carlo simulations, which show that $N_{\text{tail}} = 1760$ guarantees that $\hat{\lambda}$ is within 5% of the true value at the 95% confidence level. Note that systematic biases could also arise if, for instance, the tail of the off-target distribution is not well modeled by Eq. (8), which could happen if the frequency band is polluted by a detector artifact or other non-Gaussianities. We do not try to correct for such biases here, and instead rely on subsequent follow-up (e.g., vetoes) to eliminate outliers arising from non-Gaussian noise.

After fitting the off-target distribution for a given ASAF candidate, we calculate the false alarm probability per

³Looking ahead to Sec. VA, there are at least 11 candidates where the subband contains excessive non-Gaussian noise, causing the off-target distributions to be poorly fit by an exponential model.




FIG. 6. Tail distribution of the log-likelihood statistic, $p(\mathcal{L})$, obtained from off-target simulations around the ASAF candidate at $f_0 = 708.65625$ Hz, $\alpha_0 = 4.28$ rad, $\delta_0 = -1.11$ rad. The off-target samples are shown by the blue histogram. The solid red line is the maximum likelihood fit using an exponential function. The shaded red region represents the 2σ fit confidence band. The black dashed line is the threshold calculated using the method in Sec. IVA.

template, i.e., the probability that any single template is above-threshold, as

$$\alpha = \int_{\mathcal{L}_{\text{th}}}^{\infty} p(\mathcal{L}) d\mathcal{L} = k e^{-\hat{\lambda}(\mathcal{L}_{\text{th}} - \mathcal{L}_{\text{tail}})}. \quad (11)$$

Individual templates are required to reach a higher log-likelihood threshold due to a trials factor caused by searching multiple templates per candidate. The false alarm probability over a search involving $N_{\text{tot}} = N_{\text{sky}} N_f$ statistically independent templates is related to the single-template false alarm probability by

$$\alpha_{N_{\text{tot}}} = 1 - (1 - \alpha)^{N_{\text{tot}}}. \quad (12)$$

We adopt $\alpha_{N_{\text{tot}}} = 0.05$ as the target false alarm threshold per candidate in our analysis, meaning we should expect to accumulate $0.05 \times 515 \sim 26$ false alarms resulting from random noise across all ASAF candidates. In practice, we invert Eq. (12) to solve for α , and then compute⁴ \mathcal{L}_{th} via Eq. (11),

⁴Since the threshold depends on a fit parameter ($\hat{\lambda}$), one could ask whether the threshold itself has some uncertainty. Ultimately, the log-likelihood threshold that one uses to discriminate between “interesting” and “uninteresting” outliers is a cutoff with no uncertainty. What is uncertain is the false alarm probability that one believes their threshold corresponds to. That is, the chosen threshold corresponds to a false alarm probability close to, but not exactly the same as the target false alarm probability, $\alpha_{N_{\text{tot}}}$.

$$\mathcal{L}_{\text{th}} = -\frac{1}{\lambda} \log(\alpha/k) + \mathcal{L}_{\text{tail}}. \quad (13)$$

Given an outlier with log-likelihood \mathcal{L}^* from a search with N_{tot} templates, we can also calculate the probability of obtaining a higher log-likelihood $\mathcal{L} > \mathcal{L}^*$ due to random noise as

$$p_{\text{noise}} = 1 - [1 - k e^{-\lambda(\mathcal{L}^* - \mathcal{L}_{\text{tail}})}]^{N_{\text{tot}}}. \quad (14)$$

The quantity p_{noise} is convenient for comparing results from different ASAF candidates, as it takes into account the trials factor, N_{tot} , which differs between candidates. Note that Eq. (12) assumes the templates are uncorrelated. However, some template correlations are expected, as the template grid has been deliberately tuned to prevent large mismatches (Sec. III C). The effective number of independent templates is difficult to compute, owing to the large number of candidates to follow-up and wide parameter ranges, and is not done here. Future studies are needed to develop ways of mitigating or correcting for correlated templates in the context of HMM CW searches, e.g., by following the method of Ref. [87].

Because our targets in this search are informed by the results of the ASAF analysis, our search is not a blind search, which incurs an additional trials penalty. The problem of rigorously quantifying how each stage of a hierarchical analysis affects CW detection confidence is challenging to address either empirically (due to computational limitations) or analytically (due to each analysis employing different detection statistics). This is ultimately outside the scope of this work, so we report trials factors separately from the ASAF analysis. In Appendix C, we examine whether the ASAF detection statistic is correlated with \mathcal{L}^* or p_{noise} . In summary, we find no correlation with \mathcal{L}^* and a possible correlation with p_{noise} .

B. Veto procedure

We apply three separate vetoes to any outliers resulting from our search: the known lines veto, single interferometer veto, and DM-off veto. These vetoes compare the observed properties of an outlier to those expected of non-Gaussian detector artifacts, and have commonly been used in HMM CW searches [e.g., [36–38,80,83]]. Outliers which do not pass any of the vetoes are likely to be of instrumental origin and are discarded. We briefly summarize the three vetoes here.

1. Known lines veto

The sensitive band of the Advanced LIGO detectors contains numerous documented spectral artifacts of instrumental origin [78]. These artifacts can introduce excess power in the data that is not always well modeled by background/off-target estimates, resulting in spurious signals. The known lines veto [83] checks whether the

frequency path of the outlier, $f^*(t)$, overlaps with any of the vetted⁵ lines compiled in Ref. [88], after accounting for Doppler line broadening due to the Earth’s motion. Specifically, the outlier is vetoed if its frequency path satisfies

$$|f^*(t_n) - f_{\text{line}}| < \frac{v_{\oplus}}{c} f_{\text{line}}, \quad (15)$$

for any time t_n , where f_{line} is the line frequency,⁶ and v_{\oplus} is the Earth’s average orbital speed.

2. Single interferometer veto

Outliers representing true astrophysical signals should be present at both detectors. Thus, repeating the search on data from just one detector should naturally decrease its detection statistic. On the other hand, if the outlier arises from an instrumental artifact affecting one of the detectors, then a single-detector analysis should increase the outlier’s detection statistic in the detector where it is present, while decreasing it in the detector where it is absent. This motivates the single interferometer veto, where the outlier template is searched at each detector individually, giving two new log-likelihoods $\mathcal{L}_a < \mathcal{L}_b$. As in Ref. [83], the outlier is vetoed if the following conditions are all true: (a) the smaller single-detector statistic is subthreshold, $\mathcal{L}_a < \mathcal{L}_{\text{th}}$, (b) the other is above the (two-detector) log-likelihood, $\mathcal{L}_b > \mathcal{L}^*$, and (c) the frequency path associated with \mathcal{L}_b overlaps with the outlier frequency path, i.e.

$$|f_b(t_n) - f^*(t_n)| < \frac{v_{\oplus}}{c} f^*(t_n), \quad (16)$$

for any t_n . If the outlier satisfies Eq. (16), then it is likely associated with noise at one of the detectors and is vetoed. Otherwise, the behavior is consistent with either a weak signal, or a noise artifact common to both detectors, and cannot be vetoed.

3. DM-off veto

Instrumental artifacts are not expected to follow the Doppler modulation patterns of real signals originating from a particular point on the sky. The DM-off veto checks how disabling sky demodulation within the \mathcal{F} -statistic computation affects the log-likelihood of the outlier [89,90]. If the log-likelihood increases, $\mathcal{L}_{\text{DM-off}} > \mathcal{L}^*$, and the recovered frequency path is still overlapping with the original, DM-on frequency path [as determined by Eq. (16)], then the outlier is likely noise and is vetoed.

⁵“Vetted” means the line is most likely nonastrophysical, whether or not an exact cause has been identified.

⁶The vetted lines also have intrinsic widths, which are sometimes asymmetric about f_{line} . These widths are simply added to the left or right bounds of Eq. (15).

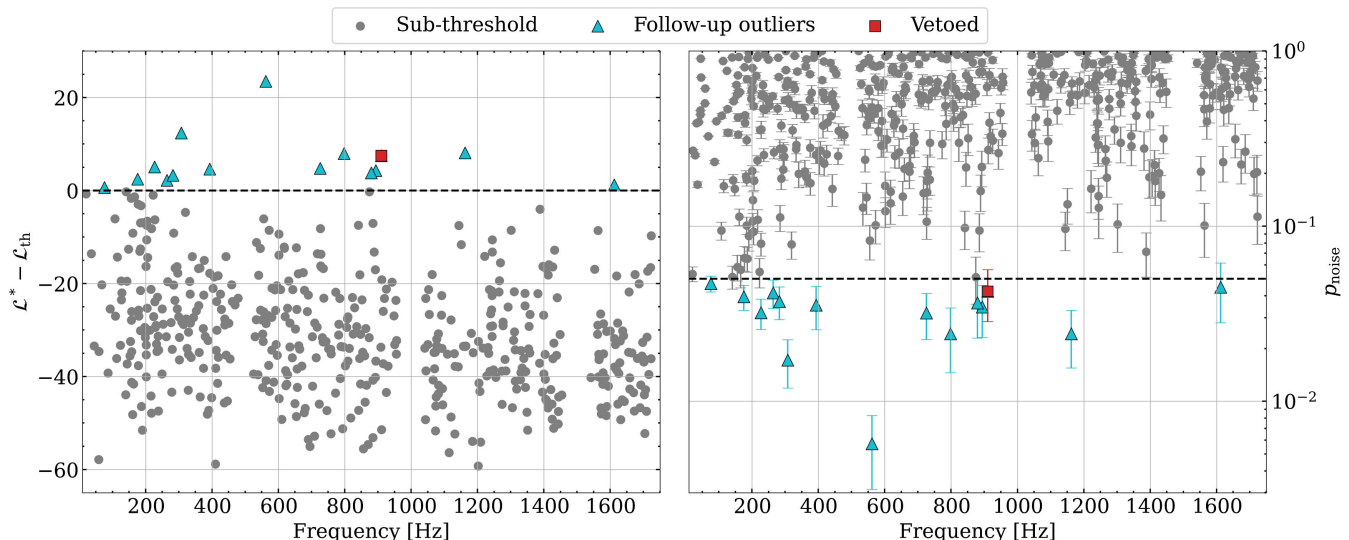


FIG. 7. Follow-up search results for the ASAF candidates. The scatter points correspond to the maximum log-likelihood template obtained for each ASAF candidate. In the left panel, we show the log-likelihood difference $\mathcal{L}^* - \mathcal{L}_{\text{th}}$ on the vertical axis, where \mathcal{L}^* is the template log-likelihood and \mathcal{L}_{th} is the threshold. The central frequency of the ASAF candidate is given on the horizontal axis. Templates which surpass the threshold, indicated by the cyan triangles, are identified as outliers and flagged for further follow-up. Sub-threshold templates are in gray. There are a total of 15 outliers, of which one (red square at ≈ 910 Hz) is vetoed. In the right panel, the vertical axis is replaced with p_{noise} for each template, with error bars indicating the 2σ statistical uncertainty propagated from the fit to the off-target distribution, i.e. they do not include systematics related to template correlations or data calibration. The dashed horizontal line is the 5% false alarm threshold. Note that the left panel does not include error bars, as log-likelihood statistics (and the thresholds) are exact quantities.

V. SEARCH RESULTS

We present the results of our follow-up of all 515 ASAF candidates in Sec. VA. Our searches yield 14 outliers that survive all vetoes. In Sec. VB, we perform deeper follow-up of these outliers by searching with a $T_{\text{coh}} = 2\text{d}$ coherence time, ultimately finding no strong evidence that these outliers represent astrophysical signals. We remind the reader that, as discussed in Sec. VA, our chosen probability of false alarm per ASAF candidate is 5%. Thus, 14 outliers from 515 candidates is statistically congruent with all of them being false alarms. We also present sensitivity estimates with respect to isolated and long-period binary sources in Sec. VC.

A. ASAF candidates

Figure 7 summarizes our search results; for each ASAF candidate, we show the maximum log-likelihood template, highlighting 15 outliers which pass the log-likelihood threshold corresponding to a 5% false alarm probability per candidate. Table I lists the outliers and their properties. All but one of these outliers pass the three vetoes listed in Sec. IV B. The single vetoed outlier terminating at 910.0801 Hz overlaps with the third harmonic of a H1 beam splitter violin mode.

Both Fig. 7 and Table I include the values of p_{noise} calculated with Eq. (14). We obtain the smallest p_{noise} value for the outlier terminating at 561.67839 Hz, for which we

estimate the false alarm probability at $p_{\text{noise}} = 0.006$. Given that we search 515 candidates, we already expect to see ~ 3 outliers at this level of significance. If we set a global threshold at 5% across all candidates, then the per-candidate false alarm probability would be $1 - (1 - 0.05)^{1/515} = 9.96 \times 10^{-5}$. As we have no outliers with p_{noise} less than this threshold, we determine that none of the outliers are statistically significant. Nonetheless, the 14 surviving outliers are subjected to further signal consistency tests in Sec. VB. The optimal frequency evolution recovered by the Viterbi algorithm is shown for each outlier in Appendix A.

Note that 11 of the subthreshold points are off-scale in Fig. 7, all of which correspond to low-frequency subbands polluted by non-Gaussian instrument noise. In these cases, the tail of the off-target distribution is poorly fit by an exponential function, and the resulting threshold is unreliable.

There are six subthreshold templates where the 2σ statistical uncertainty in p_{noise} includes the per-candidate false alarm threshold, even though they are not considered outliers by our $\mathcal{L}^* > \mathcal{L}_{\text{th}}$ cutoff. The uncertainty in p_{noise} is propagated from the uncertainty in $\hat{\lambda}$, and reflects the fact that we cannot know exactly what threshold value corresponds to the target false alarm probability $\alpha_{N_{\text{tot}}}$ based on a fit to the off-target distribution. While it is arguable that these six barely subthreshold templates should be treated as outliers, we do not consider it worthwhile to do so, owing

TABLE I. Properties of the 15 outliers recovered by our search. The first three columns pertain to the ASAF candidate: the central frequency of the 1/32 Hz subband, f_0 , and the central right ascension and declination coordinates of the sky pixel, (α_0, δ_0) . The fourth column gives the total number of $(\alpha, \delta, \dot{f})$ templates searched for that candidate, N_{tot} , and the fifth column is the log-likelihood threshold, \mathcal{L}_{th} . The next six columns give the properties of the maximum log-likelihood search template: sky position, (α, δ) ; frequency derivative, \dot{f} ; terminating frequency of the demodulated Viterbi path, $f(t_{N_T})$, Viterbi log-likelihood, \mathcal{L}^* , and the probability that a search of the ASAF candidate produces an outlier with $\mathcal{L}^* > \mathcal{L}_{\text{th}}$ from noise, p_{noise} . The last column shows whether (Y) or not (N) the outlier is vetoed. All numerical $(\alpha, \delta, \dot{f})$ values are rounded to the first significant digit of the grid spacing.

f_0 [Hz]	α_0 [rad]	δ_0 [rad]	N_{tot}	\mathcal{L}_{th}	α [rad]	δ [rad]	\dot{f} [10^{-10} Hz/s]	$f(t_{N_T})$ [Hz]	\mathcal{L}^*	p_{noise}	Vetoed
76.1875	0.88	0.48	11	2634.7	0.88	0.48	4	76.20042	2635.4	0.047	N
175.59375	0.44	0.17	33	2650.3	0.41	0.17	-6	175.58405	2652.8	0.04	N
227.0625	2.65	-0.30	55	2652.2	2.60	-0.30	-8	227.07017	2657.3	0.032	N
264.0	0.83	0.08	55	2657.9	0.80	0.08	-2	263.98836	2660.1	0.042	N
282.3125	4.81	-0.68	88	2655.7	4.79	-0.70	-6	282.30154	2659.0	0.037	N
307.65625	3.93	-1.21	286	2738.5	3.89	-1.24	0	307.66895	2750.9	0.017	N
393.28125	0.93	1.00	330	2717.1	0.99	1.00	6	393.27929	2721.8	0.035	N
561.6875	5.30	1.366	1892	2782.7	5.34	1.383	-2	561.67839	2806.2	0.006	N
725.8125	4.173	-0.43	429	2659.5	4.180	-0.46	-6	725.8251	2664.2	0.032	N
798.25	1.10	-1.315	3124	2782.0	1.22	-1.339	-10	798.26468	2789.9	0.024	N
879.75	3.75	1.107	2453	2741.1	3.77	1.091	-10	879.75653	2744.8	0.036	N
893.78125	0.246	-0.623	1023	2673.9	0.246	-0.582	8	893.78092	2678.2	0.035	N
910.09375	4.566	0.524	1243	3470.2	4.546	0.537	10	910.0801	3477.7	0.042	Y
1162.34375	3.780	-0.524	1309	2672.6	3.763	-0.533	-2	1162.3295	2680.7	0.024	N
1612.21875	2.649	-0.572	3498	2686.8	2.646	-0.563	-2	1612.21734	2688.1	0.045	N

to their doubly subthreshold status (a subthreshold outlier associated with a subthreshold ASAF candidate).

B. Outlier analysis

We further investigate the 14 post-veto surviving outliers by reanalyzing each ASAF candidate that yielded an outlier with twice the original coherence time, $T_{\text{coh}} = 2d$. In accordance with the T_{coh} scaling laws [4], we also double the number of frequency bins and increase the number of $(\alpha, \delta, \dot{f})$ templates by a factor of 16.⁷ It is worth noting that increasing T_{coh} unavoidably changes the signal model by limiting the allowed amount of spin-wandering. However, given that $T_{\text{coh}} = 1d$ is already much shorter than the expected spin-wandering timescale of known neutron stars (see Sec. III B), there is some allowance for doubling T_{coh} from an astrophysical standpoint. For each outlier, we check whether the maximum log-likelihood template from the $T_{\text{coh}} = 2d$ search is no more than one grid step away (in sky position or \dot{f}) from the loudest $T_{\text{coh}} = 1d$ template, where the grid step size corresponds to the $T_{\text{coh}} = 1d$ template spacing. None of the outliers satisfy this closeness

⁷The \dot{f} spacing goes as $\Delta\dot{f} = \Delta f / T_{\text{coh}}$, so doubling the coherence time and number of frequency bins leads to a factor of four increase in the number of \dot{f} templates. Similarly, the spacing in each sky coordinate scales inversely with T_{coh} , leading to a factor of four increase in the number of sky grid points. Thus, doubling the coherence time results in a factor of 16 more $(\alpha, \delta, \dot{f})$ templates.

criterion; for 12 of them, the sky positions differ by more than one grid step, and, for the other two, the template \dot{f} values differ.

Simply comparing the loudest templates does not conclusively rule out these outliers as noise, especially at marginal significance, as the $T_{\text{coh}} = 2d$ search could reveal a noise fluctuation which shifts the maximum likelihood to a different template. Thus, it is useful to examine the Viterbi scores, defined as [38,80]

$$\mathcal{S} = \frac{\mathcal{L}_{q^*} - \mu_{\mathcal{L}}}{\sigma_{\mathcal{L}}}, \quad (17)$$

where $\mu_{\mathcal{L}}$ and $\sigma_{\mathcal{L}}^2$ denote the log-likelihood mean and variance over all N_Q terminating paths for a single $(\alpha, \delta, \dot{f})$ template,

$$\begin{aligned} \mu_{\mathcal{L}} &= \frac{1}{N_Q} \sum_{i=1}^{N_Q} \mathcal{L}_{q_i}, \\ \sigma_{\mathcal{L}}^2 &= \frac{1}{N_Q} \sum_{i=1}^{N_Q} (\mathcal{L}_{q_i} - \mu_{\mathcal{L}})^2. \end{aligned} \quad (18)$$

Here, \mathcal{L}_{q_i} is the log-likelihood of the path ending in frequency bin q_i , and \mathcal{L}_{q^*} is the maximum log-likelihood path [Eq. (4)]. Unlike the log-likelihood, which depends on the number of steps N_T in the HMM, the score measures the significance of a given path relative to all other




FIG. 8. Comparison of the Viterbi scores between the $T_{\text{coh}} = 1\text{d}$ ($S_{1\text{d}}$, horizontal axis) and $T_{\text{coh}} = 2\text{d}$ ($S_{2\text{d}}$, vertical axis) searches. The dashed line marks where $S_{1\text{d}} = S_{2\text{d}}$. We show results for the 14 surviving outliers (cyan triangles), and for synthetic data containing either pure Gaussian noise (red rhombuses) or Gaussian noise with injected signals (blue squares). Orange circles mark injections which are “detected” by both searches, i.e., the maximum log-likelihood template is within one parameter grid step (assuming the $T_{\text{coh}} = 1\text{d}$ template spacing) from the injection parameters. The outliers are broadly consistent with the Gaussian noise scenario. When an injection is too weak to be detected, the scores fall into the Gaussian noise region, as shown by the blue squares which are not circled.

frequency paths in the same band, giving a normalized quantity that can be directly compared between searches with different T_{coh} . For $T_{\text{coh}} = 1\text{d}$, we record the score corresponding to the maximum log-likelihood template, $S_{1\text{d}}$, as before. For $T_{\text{coh}} = 2\text{d}$, we select the score of the maximum log-likelihood template, considering only templates which lie within one grid step from the loudest $T_{\text{coh}} = 1\text{d}$ template, where again the grid step size corresponds to the $T_{\text{coh}} = 1\text{d}$ search. If the original $T_{\text{coh}} = 1\text{d}$ template represents a real signal that is loud enough to be detected, then we expect a more sensitive search over the nearby parameter space to recover the same signal at higher significance, i.e. with larger Viterbi score.

As a baseline comparison, we perform $T_{\text{coh}} = 1\text{d}$ and $T_{\text{coh}} = 2\text{d}$ analyses on synthetic data consisting of either Gaussian noise with injected signals or pure Gaussian noise. For these synthetic data searches, we select a $1/32$ Hz band centered at $f_0 = 800$ Hz and a sky pixel with origin $(\alpha_0, \delta_0) = (4.271 \text{ rad}, -0.34 \text{ rad})$, and search it exactly as done for the ASAF candidates. We generate 200 datasets for each scenario (Gaussian noise, Gaussian noise with an injection) with different realizations of Gaussian noise and injection parameters. The injected signals are sampled similarly to the mismatch study in Sec. III C 3, i.e. the injections are randomly injected into the sky pixel and subband with \dot{f} sampled uniformly between $\pm 10^{-9}$ Hz/s.

The amplitudes, h_0 , are sampled uniformly around the $T_{\text{coh}} = 1\text{d}$ detection limit, $\sqrt{S_n}/h_0 \in [35, 60]$, as informed by our sensitivity calculations in Sec. V C.

We show the results for our injection studies alongside the ASAF follow-up outliers in Fig. 8, where we observe how the scores tend to cluster around different $S_{2\text{d}}/S_{1\text{d}}$ ratios depending on whether a signal is detected or not. When the data contain an injected signal, the scores cluster around larger ratios compared to the pure Gaussian noise scenario, as the SNR is amplified relative to the background. We observe that the scores for the follow-up outliers cluster in the same region as the Gaussian noise simulations, suggesting their behavior is more consistent with Gaussian noise than an astrophysical signal. Taken with their low significance, and the fact that the $T_{\text{coh}} = 2\text{d}$ analyses recover different maximum log-likelihood templates, we conclude that there is no strong evidence in support of the outliers being genuine astrophysical signals.

C. Sensitivity

We derive empirical sensitivity estimates for our search by recovering simulated signals injected into O3 detector data. The sensitivity is defined as the minimum strain amplitude, $h_0^{95\%}$, at which 95% of all injected signals are recovered above the log-likelihood threshold. Another commonly used metric is the sensitivity depth, which measures the sensitivity relative to the noise level of the detectors [91,92]:

$$\mathcal{D}^{95\%} = \frac{\sqrt{S_n}}{h_0^{95\%}}, \quad (19)$$

where S_n is the inverse-squared average, running-median PSD for both detectors. We estimate both $h_0^{95\%}$ and $\mathcal{D}^{95\%}$ at 100 ASAF candidates that contain no outliers or instrument lines. This lets us reuse the previously calculated thresholds for those candidates, without the need for additional off-target simulations.

1. Injection campaign

We characterize our sensitivity against two types of signals: isolated sources, and sources in long-period binary systems with orbital period $P_b > 1$ yr. For each candidate and source type, we generate $N_{\text{inj}} = 800$ injections within the candidate sky pixel and subband. We further stipulate that the signal frequency must be contained within the $1/32$ Hz band. For the binary injections, we thus fix the signal frequency (at the reference time) to the central frequency of the subband, and apply a crude Doppler amplitude cut of $2\Delta f_{\text{orb}} < 1/32$ Hz, where

$$\Delta f_{\text{orb}} = \frac{2\pi a_p}{P_b} \left(\frac{1+e}{1-e} \right)^{1/2}, \quad (20)$$

is the maximum possible Doppler shift caused by the binary orbital motion of the source. This ensures the frequency path of the injected signal does not drift outside the subband. In Eq. (20), a_p is the projected semimajor axis along the line-of-sight in units of light-seconds (ls), and e is the orbital eccentricity. The remaining parameters are drawn uniformly in the same manner for both source types: frequency derivative, $\dot{f} \in [-10^{-9} \text{ Hz/s}, 10^{-9} \text{ Hz/s}]$, strain amplitude, $h_0/10^{-25} \in [0.2, 4]$, cosine inclination angle, $\cos i \in [-1, 1]$, polarization angle, $\psi \in [-\pi/2, \pi/2]$, and reference phase, $\phi_0 \in [0, 2\pi]$. For the binary injections, we also sample log-uniformly over the projected semimajor axis, $a_p \in [10 \text{ ls}, 10^5 \text{ ls}]$, binary orbital period, $P_b \in [1 \text{ yr}, 10^3 \text{ yr}]$, eccentricity, $e \in [10^{-7}, 0.95]$, and uniformly over the time of passage of the ascending node (in GPS time), $t_{\text{asc}} \in [t_0 - 10^8 \text{ s}, t_0 + 10^8 \text{ s}]$, where $t_0 = 1253764818.0$ is the midpoint GPS time of O3. We run our search algorithm on each injection, searching only the templates which are within one grid step in sky position and \dot{f} from the injection parameters. We do not search over any binary parameters, instead letting the HMM and Viterbi algorithm accommodate the frequency variation induced by the binary motion.

2. Sensitivity estimates

The fraction of detected signals at a given strain amplitude, known as the detection efficiency, $\mathcal{E}(h_0)$, is modeled as a logistic function,

$$\mathcal{E}(h_0) = \frac{1}{1 + e^{-\beta_0 - \beta_1 h_0}}, \quad (21)$$

where $\beta_{0,1}$ are fit parameters. We seek the amplitude $h_0^{95\%}$ such that $\mathcal{E}(h_0^{95\%}) = 0.95$. First, Eq. (21) must be fitted using the injection search results. This can be done by either binning the injections by h_0 and counting the number of detections in each bin, or using binary logistic regression, which avoids the need for binning. We opt for the latter approach, finding the fit parameters which maximize the log-likelihood function:

$$\ell(\beta_0, \beta_1) = \sum_{i=1}^{N_{\text{inj}}} [d_i \log(\mathcal{E}_i) + (1 - d_i) \log(1 - \mathcal{E}_i)], \quad (22)$$

where h_i is the strain amplitude of the i th injection, $d_i = 1$ if the injection is detected⁸ ($d_i = 0$ otherwise), and $\mathcal{E}_i = \mathcal{E}(h_i)$. We determine the best fit parameters numerically using Markov chain Monte Carlo (MCMC) sampling, and then transform the joint posterior distribution for $\beta_{0,1}$ to the posterior distribution for $h_0^{95\%}$ using Eq. (21), assuming

⁸An injection is considered “detected” if any of the templates in the vicinity of the injection, defined as being within one grid step in \dot{f} and sky position, pass the log-likelihood threshold.




FIG. 9. Example sensitivity estimate for an ASAF candidate. Each black cross corresponds to a simulated signal injected into O3 data, which is assigned a value of “1” if it was detected above threshold, and “0” otherwise. The binary data are fitted with a logistic regression, letting us estimate where the detection efficiency crosses the 95% threshold (gray line). The solid blue curve is the median posterior fit, with the 95% credible interval (CI) shown by the shaded blue region. The constraint on the sensitivity, $h_0^{95\%}$, is shown by the red square, with error bars corresponding to the 95% CI.

flat priors on $\beta_{0,1}$. Figure 9 shows an example of using this logistic regression technique to estimate $h_0^{95\%}$.

In Fig. 10, we report sensitivity limits for a representative sample of 100 ASAF candidates, marginalized over all other source parameters. The highest sensitivity is achieved around 222.6 Hz, where the minimum detectable strain amplitude is $h_0^{95\%} = 8.8 \times 10^{-26}$ for isolated sources and $h_0^{95\%} = 9.4 \times 10^{-26}$ for long-period binaries, corresponding to sensitive depths of 55 and 51, respectively. We observe that the depth decreases slightly at higher frequencies. Since the number of templates scales with the square of the frequency, the thresholds are higher due to the trials penalty, which reduces the sensitivity of the search. The sensitivity to long-period binaries is consistently below that of isolated sources. This is not unexpected given that we only considered isolated signals when calibrating the mismatch in Sec. III C 3. Overall, our limits are competitive with previous directed HMM CW searches which used comparable coherence times [34], as well as the minimum upper limits set by O3 all-sky CW searches [44,45] and the ASAF analysis [48]. However, we caution that these sensitivities/upper limits are conditional on the assumed signal model and false alarm threshold, and are not directly comparable with each other.

We note that the long-period binary sensitivities reported in this work rely on the signal remaining mostly contained inside the frequency search band. If the amplitude cut [Eq. (20)] or frequency centering assumptions are relaxed, our sensitivity worsens by a factor between 2–3, as the binary motion-induced Doppler modulation can easily shift the signal out of band for CW frequencies $\gtrsim 100$ Hz. Future studies interested in optimizing Viterbi for use in




FIG. 10. Sensitivity limits for our search, estimated for 100 of the 515 total ASAF candidates. The color indicates the sensitivity to isolated sources (red squares) or long-period binaries (cyan rhombuses). In the top panel, we show the minimum strain amplitude, $h_0^{95\%}$, at which we detect 95% of the injections, marginalized over all other source parameters. The horizontal axis gives the central frequency of the candidate. The black curve shows an O3 run-averaged Advanced LIGO noise spectrum, which has been scaled to match the general shape of the sensitivities. In the bottom panel, we show the sensitivity depth for the same set of ASAF candidates. The error bars on the estimated sensitivities correspond to the 95% credible interval from the logistic regression.

long-period binary searches should consider using larger frequency bands to accommodate the binary Doppler shift, and/or include long-period binary templates in the matched filter search.

VI. CONCLUSIONS

In this study, we performed a triggered CW search by following up subthreshold stochastic GW candidates. Our CW search uses a HMM framework, combined with the \mathcal{F} -statistic, to achieve a flexible signal model that is sensitive to a wide range of frequency evolution [34–38,60–66]. Using Advanced LIGO data from O3, we have applied this method to follow-up all 515 subthreshold candidates identified in the LVK O1–O3 ASAF analysis [48], marking the first time a CW search has targeted stochastic candidates. Our search used a coherence time of 24 hr and 10^{-5} Hz frequency bins, providing robustness to signals whose frequency changes by at most 10^{-5} Hz per day after

correcting for secular spin-down/up, which is enough to capture the spin-wandering behavior of most known neutron stars (e.g., Scorpius X-1 [36]).

Of the 515 ASAF candidates, we obtain 15 outliers of marginal significance passing a threshold corresponding to a 5% false alarm probability per candidate. Applying a set of vetoes eliminates one outlier, leaving 14 surviving outliers. We then subjected these outliers to additional follow-up by repeating the search with double the coherence time (48 hr). None of our outliers are recovered by this more sensitive, longer-coherence search, nor do they exhibit the increase in log-likelihood expected of CW signals, suggesting the outliers are more consistent with random noise. These investigations lead us to conclude that our outliers are unlikely to represent true astrophysical signals. Data from future observing runs may reveal whether any of the outliers persist, or disappear into the noise background.

The sensitivity of our search to isolated and long-period binary sources is investigated with injection studies. We obtain sensitivity limits comparable with current all-sky CW upper limits and the ASAF upper limits. The minimum detectable strain for isolated sources with no spin-wandering is 8.8×10^{-26} , achieved at a frequency of 222.6 Hz, and the corresponding sensitivity to long-period binaries with orbital periods greater than 1 yr is 9.4×10^{-26} . The HMM approach allows one to search for long-period binaries without having to search over any binary orbital templates, reducing the required number of search templates and associated computational cost.

Our search method has broad applicability beyond following up stochastic candidates, and could be used to follow up high-priority candidates identified by all-sky CW searches. As we assume an alternative, spin-wandering signal model, our method should prove generally flexible to follow up outliers with poor signal phase constraints, and is complementary to currently existing, fully coherent follow-up procedures which assume a constant or secularly evolving signal frequency.

ACKNOWLEDGMENTS

We thank James Clark for assistance with submitting computing jobs through the IGWN Computing Grid. We thank the LVK CW and Anisotropic working groups for helpful discussions. This material is based upon work supported by NSF’s LIGO Laboratory which is a major facility fully funded by the National Science Foundation. This research was done using services provided by the OSG Consortium [93–96], which is supported by the National Science Foundation Awards No. OAC-2030508 and No. OAC-1836650. This research was enabled in part by support provided by the Digital Research Alliance of Canada ([97]). We are also grateful for computational resources provided by the LIGO Laboratory and supported by National Science Foundation Awards No. PHY-0757058 and No. PHY-0823459. A. M. K., H. D., E. G., and J. M.

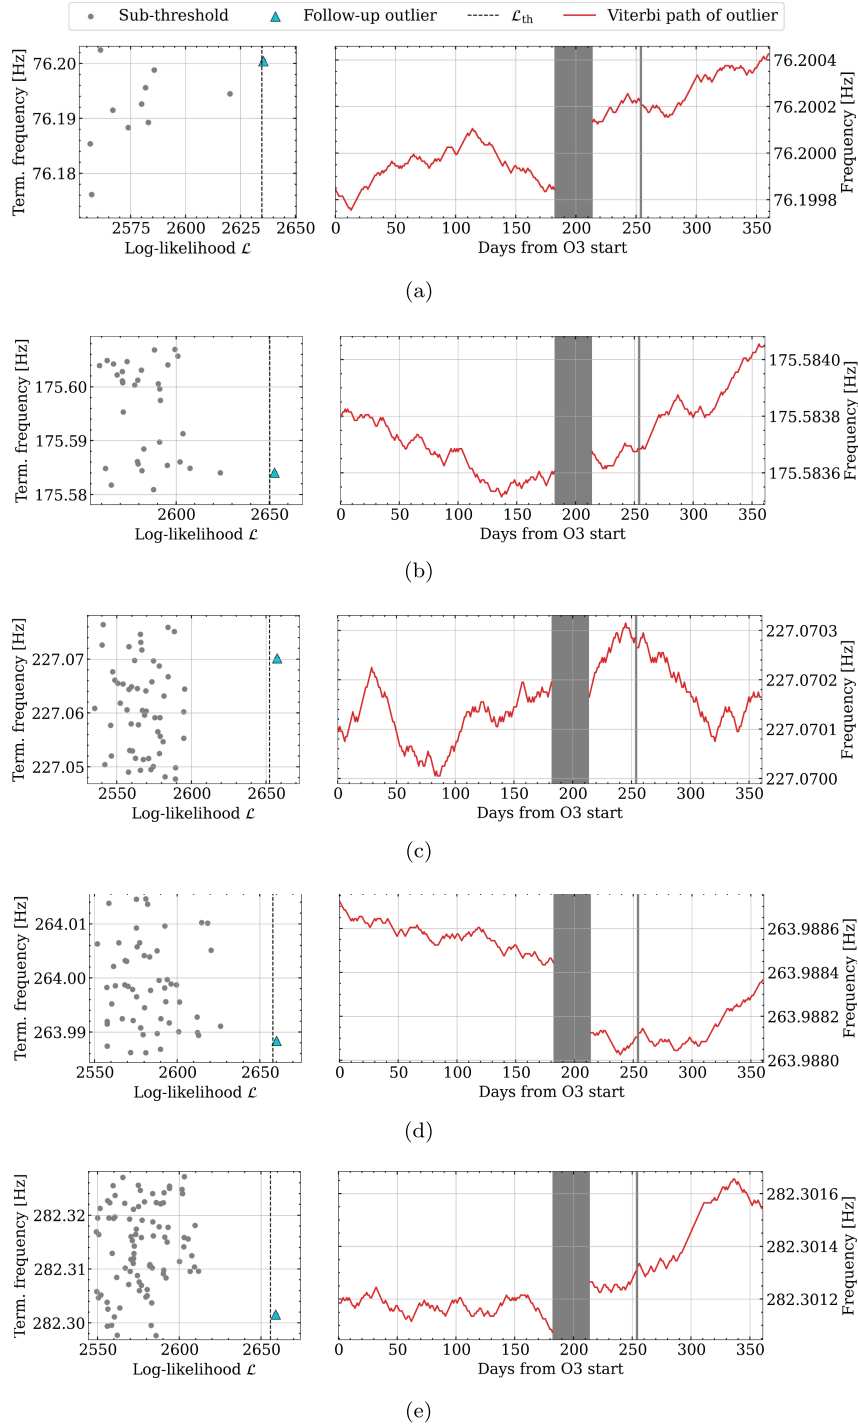


FIG. 11. Detailed search results for the 15 ASAF candidates which yielded an outlier. Left panels show the log-likelihood statistics for every searched template, with the threshold indicated by the vertical dashed line. Each template is associated with a frequency path terminating at a specific bin (after template demodulation), which is given on the vertical axis. Subthreshold templates are shown by the gray points, and the above-threshold outlier is marked as a cyan triangle. Right panels show the optimal frequency path of the outlier template. The gray vertical bands indicate where no data is available; the large gap near the middle is the O3 commissioning break. The parameters of the outlier template, $(\alpha, \delta, \dot{f})$, are stated in the caption of each subfigure, with sky positions in units of radians. Templates (a) $\alpha=0.88, \delta=0.48, \dot{f}=4 \times 10^{-10}$ Hz/s, (b) $\alpha=0.41, \delta=0.17, \dot{f}=-6 \times 10^{-10}$ Hz/s, (c) $\alpha=2.60, \delta=-0.30, \dot{f}=-8 \times 10^{-10}$ Hz/s, (d) $\alpha=0.80, \delta=0.08, \dot{f}=-2 \times 10^{-10}$ Hz/s, (e) $\alpha=4.79, \delta=-0.70, \dot{f}=-6 \times 10^{-10}$ Hz/s, (f) $\alpha=3.89, \delta=-1.24, \dot{f}=0 \times$ Hz/s, (g) $\alpha=0.99, \delta=1.00, \dot{f}=6 \times 10^{-10}$ Hz/s, (h) $\alpha=5.34, \delta=1.383, \dot{f}=-2 \times 10^{-10}$ Hz/s, (i) $\alpha=4.180, \delta=-0.46, \dot{f}=-6 \times 10^{-10}$ Hz/s, (j) $\alpha=1.22, \delta=-1.339, \dot{f}=-1 \times 10^{-9}$ Hz/s, (k) $\alpha=3.77, \delta=1.091, \dot{f}=-1 \times 10^{-9}$ Hz/s, (l) $\alpha=0.246, \delta=-0.582, \dot{f}=8 \times 10^{-10}$ Hz/s, (m) (Vetoed) $\alpha=4.546, \delta=0.537, \dot{f}=1 \times 10^{-9}$ Hz/s, (n) $\alpha=3.763, \delta=-0.533, \dot{f}=-2 \times 10^{-10}$ Hz/s, (o) $\alpha=2.646, \delta=-0.563, \dot{f}=-2 \times 10^{-10}$ Hz/s.

acknowledge funding support from the Natural Sciences and Engineering Research Council of Canada. A. M. K. acknowledges funding support from a Killam Doctoral Scholarship. J. B. C., L. S., L. D., L. S., and A. M. acknowledge the support of the Australian Research Council Centre of Excellence for Gravitational Wave Discovery (OzGrav), Project No. CE170100004. H. M. acknowledges the support of the UK Space Agency Grants No. ST/V002813/1 and No. ST/X002071/1.

APPENDIX A: OUTLIER SEARCH RESULTS

More detailed search results for the 15 outliers, including the Viterbi frequency paths, are reported in Fig. 11. Note that the outlier in Fig. 11(m) is vetoed. Visually inspecting the Viterbi paths, the majority of the outliers exhibit large jumps in frequency before and after the O3 commissioning break [Figs. 11(a), 11(d)–11(j), 11(k)–11(m), and 11(o)]. The sizes of these jumps ranges between 25–30 frequency

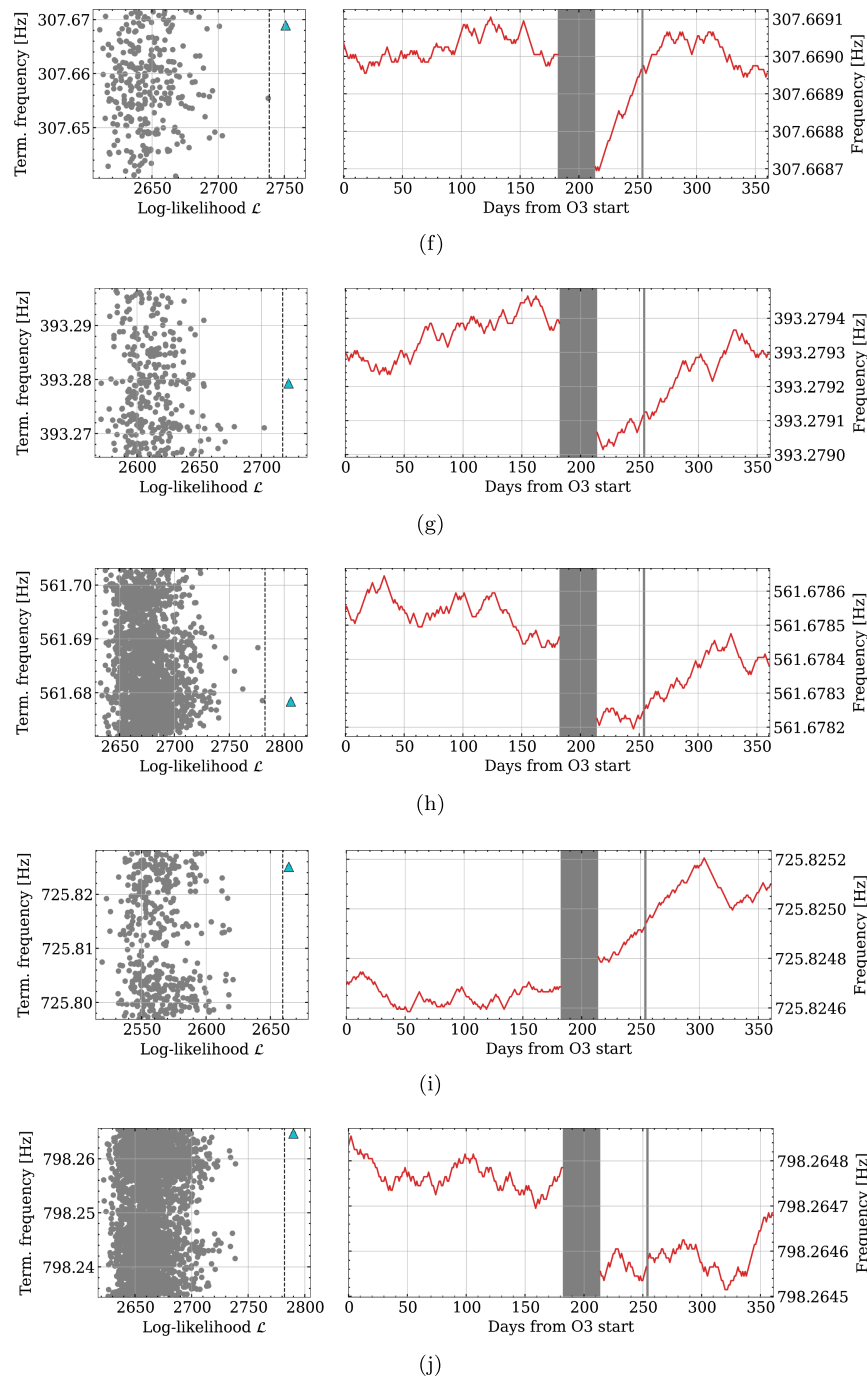


FIG. 11. (Continued)

bins, implying the frequency track moved almost exclusively upward or downward for many consecutive days. While a frequency jump of this nature is technically allowed by the random walk signal model, the coincidence with the O3 break is certainly suspicious. We ultimately discard all of these outliers based on our investigations in Sec. VB.

APPENDIX B: NUMBER OF OUTLIERS

Assuming all 14 surviving outliers represent false alarms, we obtain fewer false alarms from our search than expected, given we targeted a 5% false alarm probability per-candidate. There are a variety of possible causes for a lack of false alarms, such as a large fluctuation of the background, template correlations (which would result in

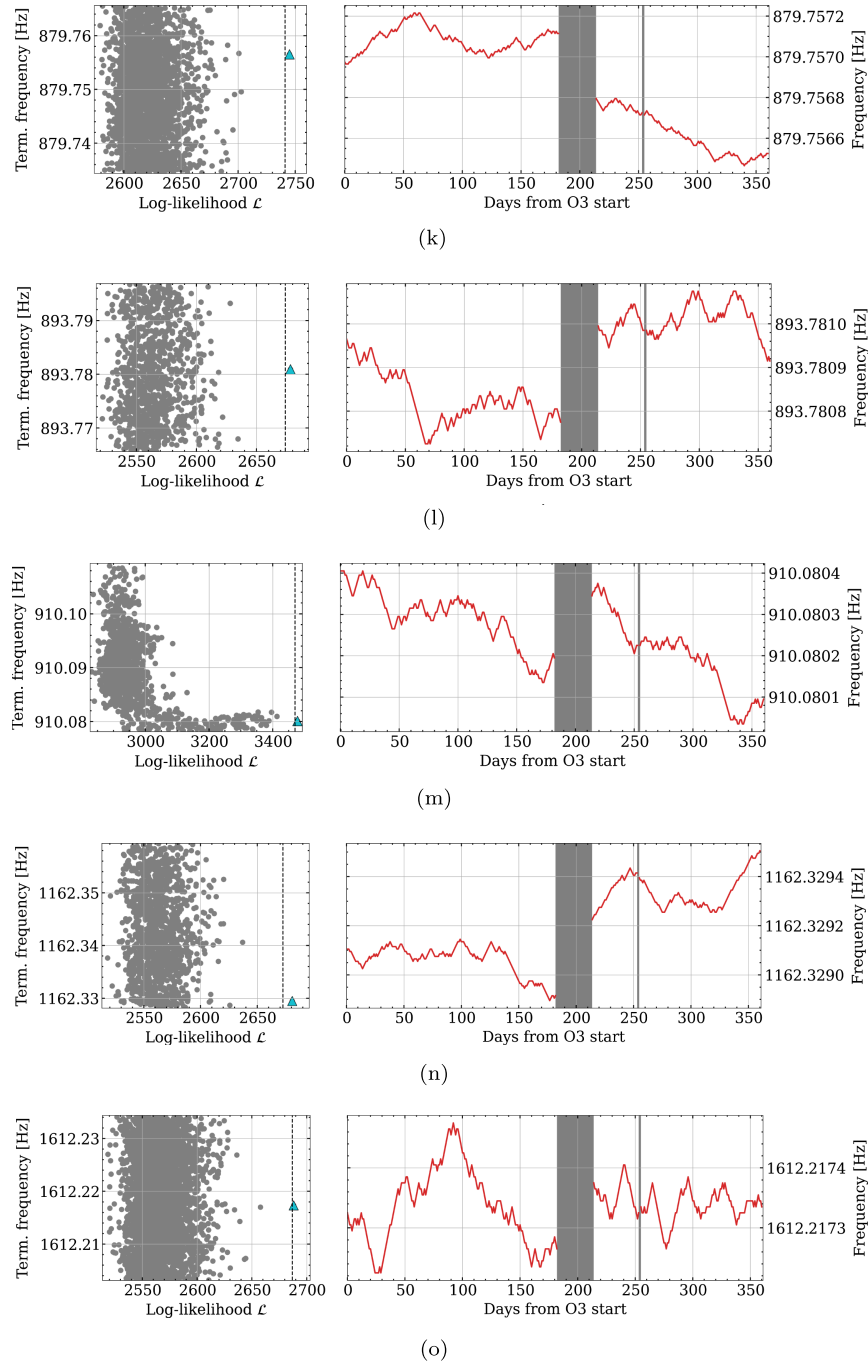


FIG. 11. (Continued)

an overestimate of the trials factor), or inaccuracy or bias in the fits to the off-target distributions. Regarding the fitting, we remind that reader that since we are following up ASAF candidates, we are in effect searching “hotspots” in the data, which may be biased toward containing unusual, non-Gaussian features. We could attempt calculating the thresholds on simulated Gaussian noise instead of using detector data, thereby avoiding complications from non-Gaussian data. Unfortunately, this is not ideal either since real detector noise is not perfectly stationary and Gaussian, and thresholds based on Gaussian noise would not reflect the true character of the data. As mentioned at the beginning of Sec. IV A, nonparametric methods for threshold estimation relax the assumption of an exponential tail, but are computationally unfeasible in a search such as ours.

Given that 14 outliers is roughly half the mean expected number of false alarms, we might estimate that our p_{noise} values are a factor of two larger than they should be. Even if we halved p_{noise} for all outliers, none would be below a global false alarm probability corresponding to 5% across all 515 candidates, and therefore our earlier conclusion that no outliers are statistically significant remains unchanged.

APPENDIX C: CORRELATION WITH ASAF SNR

Figure 12 shows p_{noise} plotted against the ASAF SNR statistic for each ASAF candidate. To test for any correlation, we compute the Pearson and Spearman’s correlation coefficients between p_{noise} and ASAF SNR, and between \mathcal{L}^* and ASAF SNR, for every ASAF candidate. For the \mathcal{L}^* and ASAF SNR comparison, the Pearson coefficient is 0.007 with (two-sided) p -value 0.877, and the Spearman’s coefficient is 0.027 with p -value 0.542, indicating no

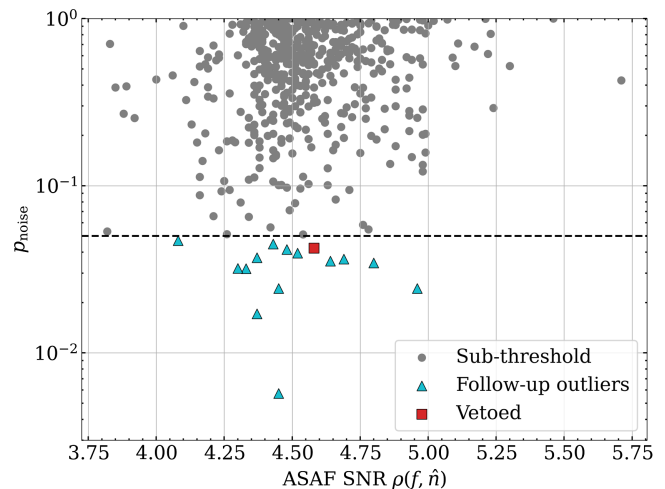


FIG. 12. Relation between p_{noise} (vertical axis), defined in Eq. (14), and the ASAF SNR (horizontal axis). Subthreshold results are shown in gray, outliers are shown by the cyan triangles, and the red square is the vetoed outlier. The horizontal dashed line corresponds to a 5% false alarm probability.

evidence for a linear or monotonic correlation between the variates. For p_{noise} and ASAF SNR, the Pearson coefficient is 0.128 with p -value 0.004, and the Spearman’s coefficient of 0.119 with p -value 0.007, suggesting that higher p_{noise} is positively correlated with higher ASAF SNR. This result is nonintuitive, as it suggests higher significance in the ASAF analysis is correlated with lower significance in our CW search, after accounting for trials factors in the CW search. Future studies may seek to investigate the correlation between CW and ASAF detection statistics when CW signals are injected upstream of both the ASAF and CW analyses.

-
- [1] P.D. Lasky, Gravitational waves from neutron stars: A review, *Publ. Astron. Soc. Aust.* **32**, e034 (2015).
 - [2] M. Sieniawska and M. Bejger, Continuous gravitational waves from neutron stars: Current status and prospects, *Universe* **5**, 217 (2019).
 - [3] B. Haskell and K. Schwenzer, Isolated neutron stars, *arXiv:2104.03137*.
 - [4] K. Riles, Searches for continuous-wave gravitational radiation, *Living Rev. Relativity* **26**, 3 (2023).
 - [5] G. Baym and D. Pines, Neutron starquakes and pulsar speedup, *Ann. Phys. (N.Y.)* **66**, 816 (1971).
 - [6] L. Bildsten, Gravitational radiation and rotation of accreting neutron stars, *Astrophys. J. Lett.* **501**, L89 (1998).
 - [7] G. Ushomirsky, C. Cutler, and L. Bildsten, Deformations of accreting neutron star crusts and gravitational wave emission, *Mon. Not. R. Astron. Soc.* **319**, 902 (2000).
 - [8] N. K. Johnson-McDaniel and B. J. Owen, Maximum elastic deformations of relativistic stars, *Phys. Rev. D* **88**, 044004 (2013).
 - [9] C. Cutler, Gravitational waves from neutron stars with large toroidal b fields, *Phys. Rev. D* **66**, 084025 (2002).
 - [10] A. Mastrano, A. Melatos, A. Reisenegger, and T. Akgün, Gravitational wave emission from a magnetically deformed non-barotropic neutron star, *Mon. Not. R. Astron. Soc.* **417**, 2288 (2011).
 - [11] P.D. Lasky and A. Melatos, Tilted torus magnetic fields in neutron stars and their gravitational wave signatures, *Phys. Rev. D* **88**, 103005 (2013).
 - [12] J. Middleton, F. E. Marshall, Q. D. Wang, E. V. Gotthelf, and W. Zhang, Predicting the starquakes in PSR J0537-6910, *Astrophys. J.* **652**, 1531 (2006).

- [13] A. I. Chugunov and C. J. Horowitz, Breaking stress of neutron star crust, *Mon. Not. R. Astron. Soc.* **407**, L54 (2010).
- [14] E. Giliberti and G. Cambiotti, Starquakes in millisecond pulsars and gravitational waves emission, *Mon. Not. R. Astron. Soc.* **511**, 3365 (2022).
- [15] A. D. Kerin and A. Melatos, Mountain formation by repeated, inhomogeneous crustal failure in a neutron star, *Mon. Not. R. Astron. Soc.* **514**, 1628 (2022).
- [16] B. J. Owen, L. Lindblom, C. Cutler, B. F. Schutz, A. Vecchio, and N. Andersson, Gravitational waves from hot young rapidly rotating neutron stars, *Phys. Rev. D* **58**, 084020 (1998).
- [17] N. Andersson, K. D. Kokkotas, and N. Stergioulas, On the relevance of the R-mode instability for accreting neutron stars and white dwarfs, *Astrophys. J.* **516**, 307 (1999).
- [18] P. Arras, E. E. Flanagan, S. M. Morsink, A. K. Schenk, S. A. Teukolsky, and I. Wasserman, Saturation of the r-mode instability, *Astrophys. J.* **591**, 1129 (2003).
- [19] C. A. van Eysden and A. Melatos, Gravitational radiation from pulsar glitches, *Classical Quantum Gravity* **25**, 225020 (2008).
- [20] M. F. Bennett, C. A. van Eysden, and A. Melatos, Continuous-wave gravitational radiation from pulsar glitch recovery, *Mon. Not. R. Astron. Soc.* **409**, 1705 (2010).
- [21] J. Aasi *et al.* (LIGO Scientific Collaboration), Advanced LIGO, *Classical Quantum Gravity* **32**, 074001 (2015).
- [22] F. Acernese *et al.* (VIRGO Collaboration), Advanced Virgo: A second-generation interferometric gravitational wave detector, *Classical Quantum Gravity* **32**, 024001 (2015).
- [23] T. Akutsu *et al.* (KAGRA Collaboration), KAGRA: 2.5 Generation interferometric gravitational wave detector, *Nat. Astron.* **3**, 35 (2019).
- [24] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), First search for gravitational waves from known pulsars with Advanced LIGO, *Astrophys. J.* **839**, 12 (2017); *Astrophys. J.* **851**, 71(E) (2017).
- [25] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), First narrow-band search for continuous gravitational waves from known pulsars in advanced detector data, *Phys. Rev. D* **96**, 122006 (2017); *Phys. Rev. D* **97**, 129903(E) (2018).
- [26] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Searches for gravitational waves from known pulsars at two harmonics in 2015–2017 LIGO data, *Astrophys. J.* **879**, 10 (2019); *Astrophys. J.* **899**, 170(E) (2020).
- [27] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Narrow-band search for gravitational waves from known pulsars using the second LIGO observing run, *Phys. Rev. D* **99**, 122002 (2019).
- [28] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Gravitational-wave constraints on the equatorial ellipticity of millisecond pulsars, *Astrophys. J. Lett.* **902**, L21 (2020).
- [29] R. Abbott *et al.* (LIGO Scientific, VIRGO, and KAGRA Collaborations), Searches for gravitational waves from known pulsars at two harmonics in the second and third LIGO-Virgo observing runs, *Astrophys. J.* **935**, 1 (2022).
- [30] R. Abbott *et al.* (LIGO Scientific, KAGRA, and VIRGO Collaborations), Narrowband searches for continuous and long-duration transient gravitational waves from known pulsars in the LIGO-Virgo third observing run, *Astrophys. J.* **932**, 133 (2022).
- [31] L. Fesik and M. A. Papa, First search for r-mode gravitational waves from PSR J0537–6910, *Astrophys. J.* **895**, 11 (2020); *Astrophys. J.* **897**, 185(E) (2020).
- [32] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Searches for continuous gravitational waves from 15 supernova remnants and fomalhaut b with Advanced LIGO, *Astrophys. J.* **875**, 122 (2019); *Astrophys. J.* **918**, 91(E) (2021).
- [33] M. A. Papa, J. Ming, E. V. Gotthelf, B. Allen, R. Prix, V. Dergachev, H.-B. Eggenstein, A. Singh, and S. J. Zhu, Search for continuous gravitational waves from the central compact objects in supernova remnants Cassiopeia A, Vela Jr., and G347.3–0.5, *Astrophys. J.* **897**, 22 (2020).
- [34] R. Abbott *et al.* (LIGO Scientific, VIRGO, KAGRA, and Virgo Collaborations), 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).
- [35] 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).
- [36] R. Abbott *et al.* (KAGRA, VIRGO, and LIGO Scientific Collaborations), Search for gravitational waves from Scorpius X-1 with a hidden Markov model in O3 LIGO data, *Phys. Rev. D* **106**, 062002 (2022).
- [37] R. Abbott *et al.* (LIGO Scientific, VIRGO, and KAGRA Collaborations), Search for continuous gravitational waves from 20 accreting millisecond x-ray pulsars in O3 LIGO data, *Phys. Rev. D* **105**, 022002 (2021).
- [38] 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).
- [39] O. J. Piccinni, P. Astone, S. D’Antonio, S. Frasca, G. Intini, I. La Rosa, P. Leaci, S. Mastrogiovanni, A. Miller, and C. Palomba, Directed search for continuous gravitational-wave signals from the Galactic center in the Advanced LIGO second observing run, *Phys. Rev. D* **101**, 082004 (2020).
- [40] R. Abbott *et al.* (KAGRA, LIGO Scientific, and VIRGO Collaborations), Search for continuous gravitational wave emission from the Milky Way center in O3 LIGO-Virgo data, *Phys. Rev. D* **106**, 042003 (2022).
- [41] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), First low-frequency Einstein@Home all-sky search for continuous gravitational waves in Advanced LIGO data, *Phys. Rev. D* **96**, 122004 (2017).
- [42] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO O2 data, *Phys. Rev. D* **100**, 024004 (2019).
- [43] R. Abbott *et al.* (KAGRA, VIRGO, and LIGO Scientific Collaborations), All-sky search for continuous gravitational waves from isolated neutron stars in the early O3 LIGO data, *Phys. Rev. D* **104**, 082004 (2021).
- [44] R. Abbott *et al.* (KAGRA, LIGO Scientific, and VIRGO Collaborations), All-sky search for continuous gravitational

- waves from isolated neutron stars using Advanced LIGO and Advanced Virgo O3 data, *Phys. Rev. D* **106**, 102008 (2022).
- [45] B. Steltner, M. A. Papa, H. B. Eggenstein, R. Prix, M. Bensch, and B. Machenschalk, Deep Einstein@Home all-sky search for continuous gravitational waves in LIGO O3 public data, *Astrophys. J.* **952**, 55 (2023).
- [46] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), 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).
- [47] P. B. Covas and A. M. Sintes, First all-sky search for continuous gravitational-wave signals from unknown neutron stars in binary systems using Advanced LIGO data, *Phys. Rev. Lett.* **124**, 191102 (2020).
- [48] R. Abbott *et al.* (KAGRA, Virgo, and LIGO Scientific Collaborations), All-sky, all-frequency directional search for persistent gravitational waves from Advanced LIGO's and Advanced Virgo's first three observing runs, *Phys. Rev. D* **105**, 122001 (2022).
- [49] LIGO, Virgo, and KAGRA Collaborations, Data for All-sky, all-frequency directional search for persistent gravitational-waves from Advanced LIGO's and Advanced Virgo's first three observing runs, Report No. LIGO-G2102029-v6, 2021, <https://dcc.ligo.org/G2102029/public>.
- [50] 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).
- [51] R. M. Shannon and J. M. Cordes, Assessing the role of spin noise in the precision timing of millisecond pulsars, *Astrophys. J.* **725**, 1607 (2010).
- [52] S. Price, B. Link, S. N. Shore, and D. J. Nice, Time-correlated structure in spin fluctuations in pulsars, *Mon. Not. R. Astron. Soc.* **426**, 2507 (2012).
- [53] A. Parthasarathy *et al.*, Timing of young radio pulsars—I. Timing noise, periodic modulation, and proper motion, *Mon. Not. R. Astron. Soc.* **489**, 3810 (2019).
- [54] N. Namkham, P. Jaroenjittichai, and S. Johnston, Diagnostics of timing noise in middle aged pulsars, *Mon. Not. R. Astron. Soc.* **487**, 5854 (2019).
- [55] M. E. Lower *et al.*, The UTMOST pulsar timing programme—II. Timing noise across the pulsar population, *Mon. Not. R. Astron. Soc.* **494**, 228 (2020).
- [56] 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).
- [57] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, The Australia telescope national facility pulsar catalogue, *Astron. J.* **129**, 1993 (2005).
- [58] P. Jaranowski, A. Krolak, and B. F. Schutz, Data analysis of gravitational-wave signals from spinning neutron stars. 1. The signal and its detection, *Phys. Rev. D* **58**, 063001 (1998).
- [59] R. Prix, Search for continuous gravitational waves: Metric of the multi-detector F-statistic, *Phys. Rev. D* **75**, 023004 (2007); *Phys. Rev. D* **75**, 069901(E) (2007).
- [60] S. Suvorova, L. Sun, A. Melatos, W. Moran, and R. J. Evans, Hidden Markov model tracking of continuous gravitational waves from a neutron star with wandering spin, *Phys. Rev. D* **93**, 123009 (2016).
- [61] 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).
- [62] J. Bayley, C. Messenger, and G. Woan, Generalized application of the Viterbi algorithm to searches for continuous gravitational-wave signals, *Phys. Rev. D* **100**, 023006 (2019).
- [63] R. Abbott *et al.* (LIGO Scientific, VIRGO, and KAGRA Collaborations), 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).
- [64] 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).
- [65] D. Beniwal, P. Clearwater, L. Dunn, L. Strang, G. Rowell, A. Melatos, and D. Ottaway, Search for continuous gravitational waves from HESS J1427-608 with a hidden Markov model, *Phys. Rev. D* **106**, 103018 (2022).
- [66] A. F. Vargas and A. Melatos, Search for continuous gravitational waves from PSR J0437-4715 with a hidden Markov model in O3 LIGO data, *Phys. Rev. D* **107**, 064062 (2023).
- [67] R. Abbott *et al.* (KAGRA, VIRGO, and LIGO Scientific Collaborations), Open data from the third observing run of LIGO, Virgo, KAGRA, and GEO, *Astrophys. J. Suppl. Ser.* **267**, 29 (2023).
- [68] J. Zweizig and K. Riles, Information on self-gating of $h(t)$ used in O3 continuous-wave searches, Report No. LIGO-T2000384-v4, 2021, <https://dcc.ligo.org/T2000384/public>.
- [69] E. Goetz and K. Riles, Segments used for creating standard SFTs in O3 data, Report No. LIGO-T2300068-v1, 2023, <https://dcc.ligo.org/T2300068/public>.
- [70] LIGO Scientific Collaboration, LIGO Algorithm Library—LALSuite, free software (GPL), 10.7935/GT1W-FZ16 (2018).
- [71] A. Buikema *et al.*, Sensitivity and performance of the Advanced LIGO detectors in the third observing run, *Phys. Rev. D* **102**, 062003 (2020).
- [72] F. Acernese *et al.* (Virgo Collaboration), Virgo detector characterization and data quality: Results from the O3 run, *Classical Quantum Gravity* **40**, 185006 (2023).
- [73] T. Akutsu *et al.* (KAGRA Collaboration), Overview of KAGRA: Calibration, detector characterization, physical environmental monitors, and the geophysics interferometer, *Prog. Theor. Exp. Phys.* **2021**, 05A102 (2021).
- [74] B. Allen, E. Goetz, D. Keitel, M. Landry, G. Mendell, R. Prix, K. Riles, and K. Wette, SFT data format version 2–3 specification, Report No. LIGO-T040164-v2, 2022, <https://dcc.ligo.org/LIGO-T040164/public>.
- [75] L. Sun *et al.*, Characterization of systematic error in Advanced LIGO calibration, *Classical Quantum Gravity* **37**, 225008 (2020).
- [76] L. Sun *et al.*, Characterization of systematic error in Advanced LIGO calibration in the second half of O3, [arXiv:2107.00129](https://arxiv.org/abs/2107.00129).
- [77] P. Covas, A. Effler, E. Goetz, P. Meyers, A. Neunzert, M. Oliver, B. Pearlstone, V. Roma, R. Schofield, V. Adya *et al.*, 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).
- [78] D. Davis, J. S. Areeda, B. K. Berger, R. Bruntz, A. Effler, R. Essick, R. Fisher, P. Godwin, E. Goetz, A. Helmling-Cornell *et al.*, LIGO detector characterization in the second and third observing runs, *Classical Quantum Gravity* **38**, 135014 (2021).
- [79] R. Prix, The \mathcal{F} -statistic and its implementation in `ComputeFstatistic_v2`, Report No. LIGO-T0900149-v6, 2019, <https://dcc.ligo.org/T0900149/public>.
- [80] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), 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); *Phys. Rev. D* **104**, 109903(E) (2021).
- [81] A. Viterbi, Error bounds for convolutional codes and an asymptotically optimum decoding algorithm, *IEEE Trans. Inf. Theory* **13**, 260 (1967).
- [82] D. Steeghs and J. Casares, The mass donor of Scorpius X-1 revealed, *Astrophys. J.* **568**, 273 (2002).
- [83] B. P. Abbott *et al.*, 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).
- [84] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere, *Astrophys. J.* **622**, 759 (2005).
- [85] A. Zonca, L. Singer, D. Lenz, M. Reinecke, C. Rosset, E. Hivon, and K. Gorski, HEALPY: Equal area pixelization and spherical harmonics transforms for data on the sphere in Python, *J. Open Source Software* **4**, 1298 (2019).
- [86] S. Suvorova, P. Clearwater, A. Melatos, L. Sun, W. Moran, and R. J. Evans, Hidden Markov model tracking of continuous gravitational waves from a binary neutron star with wandering spin. II. Binary orbital phase tracking, *Phys. Rev. D* **96**, 102006 (2017).
- [87] R. Tenorio, L. M. Modafferi, D. Keitel, and A. M. Sintes, Empirically estimating the distribution of the loudest candidate from a gravitational-wave search, *Phys. Rev. D* **105**, 044029 (2022).
- [88] E. Goetz, A. Neunzert, K. Riles, A. Matas, S. Kandhasamy, J. Tasson, C. Barschaw, H. Middleton, S. Hughey, L. Mueller, J. Heinzl, J. Carlin, A. Vargas, and I. Hollows, O3 lines and combs found in self-gated C01 data, Report No. LIGO-T2100200-v2, 2023, <https://dcc.ligo.org/LIGO-T2100200/public>.
- [89] S. J. Zhu, M. A. Papa, and S. Walsh, New veto for continuous gravitational wave searches, *Phys. Rev. D* **96**, 124007 (2017).
- [90] D. Jones, L. Sun, J. Carlin, L. Dunn, M. Millhouse, H. Middleton, P. Meyers, P. Clearwater, D. Beniwal, L. Strang, A. Vargas, and A. Melatos, Validating continuous gravitational-wave candidates from a semicoherent search using Doppler modulation and an effective point spread function, *Phys. Rev. D* **106**, 123011 (2022).
- [91] B. Behnke, M. A. Papa, and R. Prix, Postprocessing methods used in the search for continuous gravitational-wave signals from the Galactic center, *Phys. Rev. D* **91**, 064007 (2015).
- [92] C. Dreissigacker, R. Prix, and K. Wette, Fast and accurate sensitivity estimation for continuous-gravitational-wave searches, *Phys. Rev. D* **98**, 084058 (2018).
- [93] R. Pordes, D. Petravick, B. Kramer, D. Olson, M. Livny, A. Roy, P. Avery, K. Blackburn, T. Wenaus, F. Würthwein, I. Foster, R. Gardner, M. Wilde, A. Blatecky, J. McGee, and R. Quick, The open science grid, *J. Phys. Conf. Ser.* **78**, 012057 (2007).
- [94] I. Sfiligoi, D. C. Bradley, B. Holzman, P. Mhashilkar, S. Padhi, and F. Würthwein, The pilot way to grid resources using glideinWMS, in *2009 WRI World Congress on Computer Science and Information Engineering* (IEEE, 2009), Vol. 2, pp. 428–432, 10.1109/CSIE.2009.950.
- [95] OSG, OSPool (OSG, 2006), 10.21231/906P-4D78.
- [96] OSG, Open science data federation (OSG, 2015), 10.21231/0KVZ-VE57.
- [97] <http://alliance.can.ca>.