

Efficient Search for Detection Candidates Using a Peak Finder Strategy for All-Sky-All-Frequency Gravitational Wave Radiometer

Arindam Sharma,^{1,*} Deepali Agarwal,^{2,3,4,†} and Sanjit Mitra^{4,‡}

¹*Department of Physics and Astronomy, University of Mississippi, University, MS 38677, USA*

²*Department of Physics and Astronomy, University of Texas Rio Grande Valley,
One West University Boulevard, Brownsville, TX 78520, USA*

³*Centre for Cosmology, Particle Physics and Phenomenology (CP3),
Université catholique de Louvain, Louvain-la-Neuve, B-1348, Belgium*

⁴*Inter-University Centre for Astronomy and Astrophysics (IUCAA), Pune 411007, India*

The first all-sky-all-frequency (ASAF) radiometer search was conducted using data from the first three observing runs of the Advanced LIGO and Advanced Virgo detectors. The significance of this search lies in its fast and unmodeled approach, leveraging a cross-correlation technique to identify common signals across the detector network. As a result, this method serves as an excellent alternative to search for unknown or poorly modeled continuous wave sources and narrowband components of the gravitational wave (GW) background. For continuous wave sources whose waveform can be modeled, this method can serve as the first stage of a hierarchical scheme by identifying sub-threshold candidates to be followed up with more optimal but computationally expensive searches. The ASAF search, however, presently suffers from beam smearing, where multiple candidates may arise due to the same noise fluctuations, detector artifact, or a GW source. This can reduce the detection probability in follow-up analyses, especially with limited computing resources. To mitigate this issue and reduce the number of correlated and unnecessary candidates, we introduce a novel Peak Finder algorithm. This algorithm helps identifying the most representative candidates while preserving detection sensitivity, thereby allowing follow up of a much larger number of independent candidates. The reduction in correlated samples leads to *a significant reduction in False Dismissal Rate (FDR)* using the Peak Finder method compared to the Full-sky method. For instance, following up 2 Peak Finder candidates at 30 Hz reduces FDR by a factor of 3.

I. INTRODUCTION

The gravitational wave (GW) transient catalog [1] recently released by LIGO-Virgo-KAGRA (LVK) collaboration, featuring more than 200 sources, marks a stage where merging black holes and neutron stars are being routinely observed. However, aside from these transient events, persistent signals, such as continuous gravitational waves (CWs) and gravitational wave backgrounds (GWBs) are yet to be detected in the audio-frequency band [refer to the reviews [2–7] and [8–12] for sources and detection methods of CWs and GWB respectively, and references therein].

The continuous waves are everlasting quasi-periodic GWs which may be generated by spinning compact objects with deviations from axial symmetry in their distribution of mass and energy. The search sensitivity for continuous wave searches has started surpassing the spin-down limits [5] for several known pulsars, and their astrophysically motivated parameter space is being constrained via GW observations, though such signals have not been detected so far [13, 14].

GWBs are persistent signals expected to arise from the superposition of individually undetectable or unresolved astrophysical GW sources [10], as well as from

phenomena in the early universe, such as inflation, first-order phase transitions, cosmic strings [8, 11, 15]. The waveform is unpredictable for GWB due to a large number of contributing sources with unknown parameters [16], making a matched-filtering-based search infeasible. Searches for GWB rely on cross-correlating data from geographically separated detectors looking for a common excess power [12, 17]. A GWB can be categorized according to spectral and angular distributions, such as narrowband or broadband, and isotropic or anisotropic (point-like or extended sources).

Searches for anisotropic GWB rely on creating “maps” of the cross power distribution by utilizing the direction dependent time delay between detectors [18–20]. This method, known as the GW Radiometer algorithm, is essentially aperture synthesis using the Earth’s rotation. The LVK collaboration has applied this method on data up to the first part of the fourth observing run (O4a) and set upper limits on directional flux and angular power spectrum of anisotropic GWBs [21–26].

The Radiometer algorithm has the capability of detecting CW sources too [21–25, 27–29]. As an unmodeled search, this method is robust to GWs from poorly understood CW sources, such as accreting neutron stars in X-ray binaries, and exotic sources, such as boson clouds around spinning black holes. A targeted narrowband radiometer search, where the radiometer is “pointed” to a specific direction and a narrowband GW signal is searched over the sensitive frequency range for terrestrial detectors (i.e., 20–1726 Hz), has been conducted in the

* asharm17@go.olemiss.edu

† deepali.agarwal@utrgv.edu

‡ sanjit@iucaa.in

past by the LVK collaboration [21–26]. This search set constraints on GW strain from potential CW source directions such as Scorpius X-1, SN 1987A, and the Galactic Center.

The first all-sky-all-frequency (ASAF) radiometer search was conducted during the third LVK observing run (O3) [30]. This was made possible by overcoming computational constraints through the use of data folding techniques [31] and the PyStoch [32] pipeline, enabling the creation of sky maps on the HEALPix grid¹ [33, 34]. The HEALPix grid parameters were set to `nside=16`, corresponding to 3072 sky pixels and pixel size of ~ 13.4 degrees. No evidence of a point-like narrowband persistent source have been found with latest O4a results, and instead, upper-limits on the effective strain amplitude have been set [26].

The capability of the ASAF search can be further augmented by identifying detection candidates and performing follow-up analyzes using more sensitive search pipelines more optimal to search for CW sources. A total of 515 sub-threshold candidates identified by the ASAF search were followed up using Viterbi algorithm (which allows for spin-wandering) [35]. In this hierarchical approach, the search sensitivity highly depends on the number of reliable follow-up candidates.

The candidates in the ASAF search are nothing but frequency-pixel pairs. One of the challenges in this search is the large sample size, which happens to be approximately $\sim 10^8$ for `nside=16` and about 50,000 frequency bins. With the available computing resources, only a few hundred top candidates can be followed up. To reduce the loss in signal-to-noise (SNR) ratio caused by mismatches between the actual direction of a point source and the centre of the closest pixel, the sky maps are created at a sky resolution sufficiently finer than the minimum beam size at that direction. Due to fine pixels, the highly anisotropic and extended nature of the beam, the detection statistic across the pixels may not be independent. The point spread function [20] or the beam varies with signal frequency, the baseline length, and sky direction. The minimum width of the primary lobe of the beam (essentially the *diffraction limit*) for the baseline constituted by the LIGO Hanford (H) and Livingston (L) detectors has a width $\Delta\theta_{\text{Diff}}$ (in radians) [23, 36]

$$\Delta\theta_{\text{Diff}} \approx \frac{50}{f \text{ (Hz)}}, \quad (1)$$

where f is the signal frequency.

We have listed the beam sizes in Table I for a range of sample frequencies. Since the pixel resolution has to be chosen to avoid significant SNR losses across a broad frequency band, the number of pixels correlated at the lower frequencies can be very high, as shown in Table I. The number of correlated pixels is calculated using the

TABLE I: Approximate beam width (diffraction limit) and number of correlated pixels (in the diffraction limited region) for different frequencies with the LIGO Hanford-Livingston (HL) baseline. The exact numbers depend on the sky direction. In this case, we have chosen $(\text{RA}, \text{Dec}) = (6.2 \text{ h}, 45^\circ)$.

f (Hz)	30	100	200	400	700
$\Delta\theta_{\text{Diff}}$ (degrees)	95.5	28.6	14.3	7.2	4.1
# of correlated pixels	503	49	12	3	1

`query_disc` function from the `healpy` package with an example noisy sky map (without source injection). The `query_disc` function returns the pixels whose centers are within the disc defined by a vector and radius, chosen to be $(\text{RA}, \text{Dec}) = (6.2 \text{ h}, 45^\circ)$ and $\Delta\theta_{\text{Diff}}/2$, respectively. Then, we count the number of pixels inside the disc, which are mentioned in Table I.

Consequently, a large number of candidates (pixels) can be triggered due to the same cause (spurious detector disturbance or GW source). One common way to address this issue is to employ clustering in sky coordinate parameter space. The clustering algorithm may vary, but the core idea is to identify candidates arising from the same root cause, club them together, and treat them as a single entity during the follow-up stage. The clustering procedure enhances the detectability of a GW source by freeing up computational resources and allowing candidates with lower-significance to be followed up. Various clustering algorithms have also been employed in the pre-processing stage of CW search pipelines [37–41].

Here, we explore a peak-finder-based clustering statistic, where we identify the “peaks” in the sky maps for follow-up studies. A peak is defined as a pixel whose SNR exceeds that of all its nearest-neighbor pixels. The SNR values associated with these peak pixels, referred to as *Peak SNR*, are employed as a detection statistic. We test our Peak SNR statistic with simulated datasets and demonstrate the improvement in detection probability compared to the “full-sky map” statistic, where all the pixels in the sky are utilized to identify candidates for follow-up studies.

The article is structured as follows: Section II presents the formalism used in the ASAF search. In Secs. III and IV, we detail the peak finder algorithm and statistical characterization through Monte Carlo simulations, respectively. We conclude the article with a summary and a discussion on the future prospects of the methods proposed here in Sec. V.

II. ALL-SKY-ALL-FREQUENCY SEARCH

The ASAF radiometer search aims to detect persistent GW sources localized in both the sky and frequency. In this section, we first review the formalism used to characterize and search for narrowband point-like GWBs. We

¹ <http://healpix.sourceforge.net>

then discuss the similarities and differences between an anisotropic narrowband GWB and a CW source. Finally, we discuss the “full-sky map” detection statistic and its drawbacks at the end of this section.

A. Characterization of Narrowband Anisotropic GWB

The metric perturbations at a space-time point $(\vec{x}, t)$, in transverse-traceless gauge, can be expanded in terms of plane waves incoming from $\hat{\Omega}$ direction and solid angle $d^2\hat{\Omega}$, with frequency in range f and $f + df$, as [17]

$$h_{ab}(t, \vec{x}) = \sum_{A=+, \times} \int_{-\infty}^{\infty} df \int_{S^2} d^2\hat{\Omega} e_{ab}^A(\hat{\Omega}) \tilde{h}_A(f, \hat{\Omega}) e^{i2\pi f(t - \hat{\Omega} \cdot \vec{x})/c}, \quad (2)$$

where c is the speed of light. $\tilde{h}_A(f, \hat{\Omega})$ are Fourier domain strain coefficients and due to the real nature of signal $\tilde{h}_A(f, \hat{\Omega}) = \tilde{h}_A^*(-f, \hat{\Omega})$. The index A and tensor e_{ab}^A represent the state of polarization and polarization tensor of GWs.

The strain coefficients $\tilde{h}_A(f, \hat{\Omega})$ are, in principle, unknown for any kind of GW source. For instance, $\tilde{h}_A(f, \hat{\Omega})$ for a single monochromatic GW source will be proportional to $\delta(f - f_0)$, but deterministic. That is, one can design a grid of astrophysically motivated f_0 and search for a signal with frequency f_0 employing matched-filtering-based techniques, comparing the data stream to a set of templates in order to obtain high correlation. In practice, however, the parameter space to be explored is significantly larger due to uncertainties in quantities such as frequency, frequency derivatives and sky directions, etc. A “blind” fully coherent matched filtering-based search is computationally prohibitive and has been performed using sub-optimal methods, e.g., via semi-coherent and hierarchical schemes [3].

On the other hand, if there is a superposition of a large (or infinite) number of GW sources with a statistical distribution of parameter values, the strain coefficients for the GWB are no longer predictable, we need to consider an ensemble of realizations for $\tilde{h}_A(f, \hat{\Omega})$. Assuming that the GW signal is statistically stationary and unpolarized, and $\tilde{h}_A(f, \hat{\Omega})$ is normally distributed, the first two moments of $\tilde{h}_A(f, \hat{\Omega})$ can be written as [17],

$$\begin{aligned} \langle \tilde{h}_A(f, \hat{\Omega}) \rangle_h &= 0, \\ \langle \tilde{h}_A(f, \hat{\Omega}) \tilde{h}_{A'}^*(f', \hat{\Omega}') \rangle_h &= \frac{1}{4} \delta(f - f') \delta^2(\hat{\Omega} - \hat{\Omega}') \delta_{AA'} \mathcal{P}(f, \hat{\Omega}), \end{aligned} \quad (3)$$

where $\mathcal{P}(f, \hat{\Omega})$ is (polarization-averaged) one-sided power spectral density for GWB. $\mathcal{P}(f, \hat{\Omega})$ is the quantity of interest and is observable through GW experiments. The directional dependence of $\mathcal{P}(f, \hat{\Omega})$ could be well localized (into a pixel) or extended (over a large-scale structure).

In such cases, $\mathcal{P}(f, \hat{\Omega})$ could be decomposed into a suitable basis as

$$\mathcal{P}(f, \hat{\Omega}) = \sum_{\alpha} \mathcal{P}_{\alpha}(f) e_{\alpha}(\hat{\Omega}), \quad (4)$$

where $e_{\alpha}(\hat{\Omega})$ are basis vectors: 1 for isotropic source, $e_{\alpha}(\hat{\Omega}) = \delta^2(\hat{\Omega} - \hat{\Omega}_{\alpha})$ for a point source in direction $\hat{\Omega}_{\alpha}$ and $e_{\alpha}(\hat{\Omega}) = Y_{\ell m}(\hat{\Omega})$, where $Y_{\ell m}(\hat{\Omega})$ are spherical harmonic basis functions. The $\mathcal{P}_{\alpha}$ values are interpreted as the GWB power in α -th component. The primary goal of anisotropic GWB searches is to estimate $\mathcal{P}_{\alpha}(f)$, for which the usual strategy is to employ the GW radiometer algorithm [19, 20]. In this article, we will focus on mapping anisotropy of the GWB in a pixel basis. (Readers interested in the ASAF search in the spherical harmonic basis may refer to [42].)

B. Map-making with GW Radiometer

The GW radiometer algorithm uses cross-correlation of data from geographically separated observatories. It is a semi-coherent search, in which data from the full observing runs is divided into small time segments and the correlation is performed segment-wise before combining with optimal weighting factors. We briefly outline the process of constructing estimators for $\mathcal{P}_{\alpha}(f)$ using this algorithm. The time-series data from the observing runs are divided in $\tau = 192$ s long, 50% Hann-windowed overlapping time segments [26, 30]. The search begins by constructing the Cross-Spectral Density (CSD), $C^I(t, f)$, as [31, 43],

$$\mathbf{C}_f \equiv C^I(t, f) \equiv \frac{2}{\tau \overline{w_1 w_2}} \tilde{s}_{I_1}^*(t, f) \tilde{s}_{I_2}(t, f), \quad (5)$$

where t is the time stamp of a given segment, and $\overline{w_1 w_2}$ are window factors. $\tilde{s}_{I_i}(t, f)$ represents the short-term Fourier transform of data from detector I_i in the detector baseline I .

Assuming the detector noise is additive, Gaussian, and stationary within each segment, and uncorrelated with the noise of a geographically separated detector, it can be shown that [31, 44],

$$\langle C^I(t, f) \rangle_{h,n} = \sum_{\alpha} \mathcal{P}_{\alpha}(f) \gamma_{\alpha}^I(t, f), \quad (6)$$

where $\langle \cdot \rangle_{h,n}$ denotes the ensemble average over noise and GW source realizations. $\gamma_{\alpha}^I(t, f)$ is a detector and source geometry-dependent function, known as the Overlap Reduction Function (ORF), given in pixel basis by [20],

$$\begin{aligned} \gamma_f &\equiv \gamma_{\alpha}^I(t, f) \\ &= \sum_A \frac{1}{2} F_{I_1}^A(\hat{\Omega}_{\alpha}, t) F_{I_2}^A(\hat{\Omega}_{\alpha}, t) e^{i2\pi f(\hat{\Omega}_{\alpha} \cdot \Delta \vec{x}_I(t))/c}, \end{aligned} \quad (7)$$

where $\Delta\vec{x}_I(t) = \vec{x}_{I_1}(t) - \vec{x}_{I_2}(t)$ is the separation vector between the two detectors, and $F_{I_i}^A(\hat{\Omega}, t)$ represents the detector response for detector I_i .

Given the persistent nature of the GW signal of interest, observations can be integrated over the entire observing run. This offers two advantages: (i) The SNR improves as more data is accumulated. (ii) Sky coverage increases, as detector baselines have blind spots that shift across the sky due to Earth's rotation [20].

Assuming that the CSD for each time segment is a Gaussian random variable, the Maximum Likelihood (and unbiased) estimator for observable GWB anisotropy, $\mathcal{P}_\alpha(f)$, is obtained by maximizing the joint likelihood across positive and negative frequency, different times and baselines [30],

$$\hat{\mathcal{P}}_f = \mathbf{\Gamma}_f^{-1} \cdot \mathbf{X}_f. \quad (8)$$

Here, $\mathbf{X}_f$ is the narrowband *dirty map* and $\mathbf{\Gamma}_f$ is the narrowband Fisher information matrix defined as follows [30, 32],

$$\begin{aligned} \mathbf{X}_f &= \gamma_f^\dagger \cdot \mathbf{N}_f^{-1} \cdot \mathbf{C}_f, \\ \mathbf{\Gamma}_f &= \gamma_f^\dagger \cdot \mathbf{N}_f^{-1} \cdot \gamma_f, \end{aligned} \quad (9)$$

where, matrix $\mathbf{N}_f$ is covariance matrix of the CSD [31, 43, 45]. The dirty map represents the sky map as observed through the detectors. By construction, the Fisher information matrix, $\mathbf{\Gamma}_f$, serves as the noise covariance matrix for the dirty map (in the weak signal limit) and encodes information about the pixel-to-pixel correlation. In terrestrial GWB experiments, inverting $\mathbf{\Gamma}_f$ is non-trivial due to the blind spots of the baseline [20, 46–48].

The ORF exhibits time translation symmetry with a sidereal period, which is leveraged to fold the entire observing run data (spanning around hundred live days) into a single sidereal day [31, 49]. Additionally, **PyStoch** [32] is a code that fully utilizes the folded data, enabling the calculation of maps on a personal computer. This approach significantly speeds up the processing, making it hundreds of times faster, and eliminates the need to store large intermediate data files.

The folding and **PyStoch** code were used to create narrowband maps using data upto O4a observing run [26]. The ASAF radiometer search was then performed using these maps to identify persistent sources of GWs [30].

In the next section, we will briefly discuss the relationship between narrowband GWB and continuous wave signals.

C. CW Source and GW Radiometer

Both continuous wave (CW) signals and gravitational wave background (GWB) signals are expected to be persistent over time, but they differ in the length of time coherence. The cross-correlation-based search performed here is sub-optimal for CW signals, as it does

not fully exploit the long-term coherence, unlike other cross-correlation-based searches discussed in the works of Dhurandhar et al. [50] and Whelan et al. [51]. In contrast, the optimal searches require modeling the phase evolution of the signal over time. Predicting the phase evolution for unknown CW sources, or for sources with frequent frequency drift (such as glitching neutron stars or accreting neutron stars in binaries), demands covering a vast parameter space. In such cases, the ASAF radiometer search, while sub-optimal, is a robust alternative due to its insensitivity to the signal model.

A comparison of the sensitivity of the radiometer search with matched-filtering-based searches was presented in Messenger et al. [27], where a mock data challenge for Scorpius X-1 was performed. The study estimated that the radiometer search has a similar sensitivity to TwoSpect [52] and less sensitivity than CrossCorr search [50]. At the same time, the radiometer search uses less than 1% resources than the other two searches. Additionally, a CW signal may have a polarized state depending on the Earth-source geometry (i.e., the inclination and polarization angles). The radiometer search estimates, assuming an unpolarized signal, could be biased for a polarized CW signal [26, 50].

D. Full-sky Map Detection Statistic Implemented in ASAF Search

We now discuss the previous detection statistic employed to distinguish GW signals from detector noise and identify candidates for follow-up [26, 30], along with their associated drawbacks.

In the ASAF search, sky maps are created by dividing the sky into 3072 pixels for each frequency bin with a width of 1/32 Hz bin width. Data from the 20-1726 Hz frequency range are analyzed.

To avoid the (ill-conditioned) Fisher information matrix inversion problem, only the diagonal part of the matrix was used to create sky maps for $\hat{\mathcal{P}}_f$, i.e.,

$$\hat{D}_\alpha(f) \equiv \frac{X_\alpha(f)}{\Gamma_{\alpha\alpha}(f)}, \quad (10)$$

and its uncertainty and SNR were calculated using

$$\begin{aligned} \sigma_\alpha(f) &= [\Gamma_{\alpha\alpha}(f)]^{-1/2}, \\ \hat{\rho}_\alpha(f) &= \frac{\hat{D}_\alpha(f)}{\sigma_\alpha(f)}. \end{aligned} \quad (11)$$

The SNR is expected to follow a Gaussian distribution with mean zero in the presence of only Gaussian noise [30], and a non-zero mean in the presence of a GW source.

Next, the SNR is used as the detection statistic to test the null hypothesis that the data is pure Gaussian noise. A candidate frequency-pixel pair with a SNR greater than a certain threshold is considered as an outlier and




FIG. 1: Simulated SNR sky maps at two frequencies, 30 Hz and 200 Hz, with and without a source. The maps illustrate the point spread function of pixels, showing the spatial distribution of SNR values and the effect of source injection on the observed signal. The source is injected with an SNR of 3 at the sky coordinates (RA, Dec) = (6.2h, 45°), marked by the red star. The yellow triangles highlight the “peaks” (representative of correlated pixels) identified using the peak finder algorithm introduced in Sec. III.

is selected for further follow-up. The threshold SNR is determined based on trial factors (i.e., treating each frequency and pixel pair as an independent experiment) and the global p -value (i.e., the probability of observing a certain SNR anywhere in the parameter space). If no candidate exceeds the SNR threshold set by 5% of the global p -value, the sub-threshold follow-up candidates are identified for further investigation. These candidates can then be analyzed with more sensitive search methods specifically tuned to continuous wave signals. For example, see Knee et al. [35] for the follow-up of O3 ASAF search [30] sub-threshold candidates.

To identify potential sub-threshold candidates while working with a limited computing budget, we consider the distribution of the maximum pixel SNR, $\rho_{\max}(f) \equiv \max_{\alpha}[\hat{\rho}_{\alpha}(f)]$ [30]. The entire frequency range is divided into 10 Hz bins, over which the sensitivity of the radiometer does not vary significantly. For each 10 Hz bin, a threshold is set on $\rho_{\max}(f)$ corresponding to the 99th percentile of the histogram created using random-time-shifted data (representing the null distribution). Candidates from the zero-lag run with $\rho_{\max}(f)$ above the threshold are then selected for further investigation. Hence, we select approximately 3 candidates per 10 Hz range, i.e., 3 candidates per 10^6 frequency-pixel pairs.

However, the “SNR” detection statistic is affected by the oversampling of the sky grid at low frequencies, as

discussed in the next section.

E. Pixel-to-pixel Correlation of Detection Statistic

The SNR for each frequency-pixel pair, as defined in Eq. (11), serves as the detection statistic. The SNR estimators for different directions at a given frequency can be correlated. The angular spread or leakage depends on the sky direction, signal frequency, and baseline length.

Let us consider two conditions: (i) The effect on SNR across the sky given a point-like GW source. (ii) The SNR across the sky in the case of pure Gaussian noise.

In the first case, the noise-averaged (or noise-GWB-averaged) observed SNR across the sky is proportional to the convolution of the Fisher information matrix and the true SNR sky map (see Appendix A). In the second case, the average SNR is zero, but the noise covariance of the SNR vector is proportional to the Fisher information matrix. Therefore, the Fisher information matrix encodes the correlation information among the pixels across the sky. A naive calculation performed in Appendix A shows that the elements of this matrix are proportional to the “sinc” function with arguments such as angular distance, frequency, and baseline length.

The angular spread of the SNR for two example frequencies, with and without a GW source, is shown in Fig. 1. The maps are simulated using methodology described in Sec. IV A. As noted, the pixel size used for the search is smaller than the angular distance of correlation; therefore, there are correlated pixels (see Table I for the number of correlated pixels as a function of frequency). Therefore, if there were more than one potential CW candidate with different strengths in a single frequency bin, the top candidates identified in the process described in previous section may only represent the stronger source. If the stronger candidate is due to noise, and there was a true astrophysical signal which is weaker than the noise, the latter will be missed in the process.

Moreover, due to this correlation, the effective trial factors should be less than the naive number computed by the number of frequency bins times the total number of pixels. The overestimation of the number of trial factors may lead to an increase in the False Dismissal Rate (FDR). This occurs because the SNR threshold is determined by fixing the global p -value and trial factor. If the trial factor is overestimated, the resulting SNR threshold corresponds to a lower effective False Alarm Rate (FAR). A lower FAR implies higher FDR. Additionally, if there is any non-Gaussian disturbance (such as a GW source or detector artifact), it is possible that all the frequency-pixel pairs passing the SNR threshold originate from the same frequency.

Although the method used in Abac et al. [26] to identify sub-threshold candidates is currently immune to pixel-to-pixel correlation, in the future, if the computing budget allows for following up on more candidates—such as more than one candidate per frequency bin—then it




FIG. 2: The flow chart illustrates the Peak Finder algorithm for a single frequency bin. For more details, refer to the main text in Sec. III.

will be necessary to avoid unnecessary follow-up of each correlated sample.

In this article, we propose a Peak SNR statistic to group the correlated samples together and identify their representative.

III. PEAK SNR STATISTIC

In this section, we describe the working principle of the “Peak Finder” algorithm.

The algorithm is presented in the flow chart in Fig. 2. We start with the SNR sky map for each frequency and information about the 8 neighboring pixels on the HEALPix grid. The procedure begins at the 0th pixel, and we check if its SNR is greater than that of its neighboring pixels. If it is, the pixel is added to the “Peaks” array. Regardless of the outcome, we move to the next pixel and repeat the comparison with its neighboring pixels. This continues for all pixels in the map. At the end, the identified “Peaks” are returned as the output. Examples of the identified peaks in the simulated sky maps are shown in Fig. 1.

IV. SIMULATION STUDY

We aim to understand the properties of the Peak SNR statistic and its effectiveness in enhancing the detection probability of GW sources through Monte Carlo simulations. In this section, we begin by describing the method used to simulate sky maps, followed by an analysis of the results obtained after applying the Peak Finder algo-

rithm. We present three key sets of results: (i) the reduction in the sample size (or trial factors), (ii) the empirical probability distribution of the statistic and its Receiver Operating Characteristic (ROC) curve, and (iii) the impact on the detection probability of injected sources.

A. Simulating SNR Sky Maps

As described in Sec. IIB, the dirty map $\mathbf{X}_f$ and its covariance matrix $\mathbf{\Gamma}_f$ (or Fisher information matrix) capture sufficient description of the data for our purpose, encapsulating the detector noise and GW signal information compressed from the time domain into the frequency-pixel domain. Therefore, we can leverage the noise covariance matrix $\mathbf{\Gamma}_f$ in the pixel domain to simulate both pure noise and noise with GW source sky maps.

We use noise sensitivity estimates saved as the folded data from the third LVK observing run, which are publicly available on Zenodo [53]. The folded data covers the frequency range of 20–1726 Hz with a bin width of 0.03125 Hz. Using this data, we compute the Fisher Information Matrix $\mathbf{\Gamma}_f$ using the PyStoch code, employing a sky resolution of $n_{\text{side}} = 16$.

Next, we fix the injection SNR $\rho_\alpha(f)$ and compute $\mathcal{P}_\alpha(f)$ using the following relation:

$$\mathcal{P}_\alpha(f) \equiv \rho_\alpha(f) \times \sigma_\alpha(f), \quad (12)$$

and Eq. (11). For example, if we inject a point source in only the i th pixel, then $\rho_i(f) \neq 0$, while the SNR remains zero for all other pixels. We then convolve the $\mathcal{P}_\alpha(f)$ sky map with the matrix $\mathbf{\Gamma}_f$ and add a noise map $\mathbf{n}_f$, such that

$$\begin{aligned} \mathbf{X}_f &= \mathbf{\Gamma}_f \cdot \mathcal{P}_f + \mathbf{n}_f, \\ \mathbf{n}_f &\sim \mathcal{N}[0, \mathbf{\Gamma}_f], \end{aligned} \quad (13)$$

where $\mathcal{N}[0, \mathbf{\Gamma}_f]$ represents a multivariate Gaussian distribution. Finally, the SNR map $\hat{\rho}_\alpha(f)$ is obtained using Eq. (11).

B. Results of Peak Finder

1. Reduction in False Detections

Simulating SNR sky maps using the method described in the previous section, we apply the Peak Finder algorithm to identify peaks in the sky map. The identified peaks, for example, at frequencies of 30 and 200 Hz, are shown in Fig. 1 for both noise-only and noise+injection source cases. As shown in the figure, for the noise+injection source case, the identified peak is located near the true injection position, though its location exhibits a frequency-dependent uncertainty that is larger at lower frequencies.

Next, we generate $N_{\text{sim}} = 10^4$ noise realizations and inject the same GW source to the noise maps. The GW




FIG. 3: Mean number of candidates as a function of SNR threshold. Plots are provided for both the All-sky (with dashed line) and Peak Finder (with solid line) methods across a range of frequencies (30-700 Hz). We note that there is a significant reduction in the number of candidates at lower frequencies with latter method.




FIG. 4: Number of peaks for different frequencies with the mean and standard deviation calculated over 10^4 maps

source parameters are fixed with an injection SNR of 3, a fixed sample sky location chosen as (RA, Dec) = (6.2 h, 45°). For the “Full-Sky” method, a pixel in a realization is considered a candidate if its SNR is greater than the chosen SNR threshold, between 0 - 5. For the Peak Finder method, we classify a pixel as a candidate if it is a peak and its SNR is greater than the same SNR threshold for each realization. Then the mean number of candidates is defined as $\frac{N_{\text{cand}}}{N_{\text{sim}}}$, where N_{cand} is the total number of candidates across all realizations. This is repeated for a range of sample frequencies: 30, 100, 200, 400, and 700 Hz.

Fig. 3 shows the reduction in the mean number of candidates (or trial factors) with the Peak Finder algorithm compared to the Full-Sky map method. The mean number of samples per map above a threshold is significantly reduced at lower frequencies (30–200 Hz). Specifically,

at an SNR threshold of 0, the ratio of candidates in the Full-Sky method to the Peak Finder method is as follows: (30 Hz: 782:1), (100 Hz: 133:1), (200 Hz: 40:1). At an SNR threshold of 4.75, the corresponding ratios are: (30 Hz: 78:1), (100 Hz: 10:1), (200 Hz: 3:1). However, at higher frequencies (400 and 700 Hz), the reduction factor is smaller, around (10:1 – 1:1) and (4:1 – 1:1), respectively. This is because the point spread function is frequency-dependent, and at high frequencies like 400 and 700 Hz, the pixel size is close to or smaller than the diffraction limit. Therefore, there are a much larger number of peaks at such frequencies, as shown in Fig. 4. The reduction in sample size is also apparent in sky maps shown in Fig. 1. For example, only three or two peaks are identified in total for maps at 30 Hz. This highlights the significant reduction in the number of samples when applying the Peak Finder method compared to the Full-Sky approach.

2. Receiver Operating Characteristic (ROC) curves

In the next step, we aim to characterize the probability density function (PDF) of the Peak SNR statistic. We use Monte Carlo simulations to empirically estimate the PDF of the Peak SNR statistic. Once we have the PDF, we can compute the False Alarm Rate, and False Dismissal Rate, which are key metrics for evaluating the performance of the Peak Finder algorithm in identifying GW sources or detection candidates.

Example PDFs for the noise-only and noise+injection source cases are shown in the first panel of Fig. 5. The histogram for the noise+injection case is obtained by simulating 10^4 noise realizations. The GW source parameters are fixed as described in the previous section. We then collect all the peaks that lie within the diffraction-limited region (see Eq. (1)) of the injected sky coordinates and plot the histogram of their SNRs. To obtain the noise-only histogram, we repeat the same procedure, but with noise-only maps. This allows us to compare the distribution of the Peak SNR statistic in both cases and evaluate how the presence of a GW source affects the statistical behavior of the peaks.

Now, in the next step, we study the effect of the injection SNR on the FAR and FDR. We define our null hypothesis as the data being pure Gaussian noise. The alternative hypothesis is that there is one point-like GW source present in the data along with Gaussian noise with non-zero power, i.e., $\mathcal{P}_\alpha > 0$. The FAR and FDR are then defined as follows:

$$\begin{aligned} \text{FDR} &= 1 - \frac{M_{\text{TD}}}{M_{\text{NI}}}, \\ \text{FAR} &= \frac{M_{\text{FA}}}{M_{\text{N}}}. \end{aligned} \quad (14)$$

where M_{TD} is the number of noise+injection maps with at least one peak in the diffraction-limited region above




FIG. 5: Left panel: The distribution of Peak SNR for noise-only (blue dashed) and noise+injection source (orange solid). Middle panel: Variation of FAR (solid) and FDR (dashed) as a function of the threshold on the Peak SNR for a set of frequencies. The error bars depict 1σ Poisson fluctuations in FAR and FDR. Right panel: Frequency-direction-averaged ROC curve for a set of injected SNR values. The threshold decreases from left to right. The error bars represent 1σ fluctuations across seven frequency bins and four sky directions. See main text in Sec. IV B 2 for details.

the SNR threshold and M_{NI} is number of noise+injection maps with at least one peak in the diffraction-limited region. M_{FA} is the number of noise-only maps with at least one peak in the diffraction-limited region above the SNR threshold and M_N is the number of noise only maps with at least one peak in the diffraction-limited region.

In the second column of Fig. 5, we show the variation of FDR (dashed) and FAR (solid) with the Peak SNR threshold. The curves are generated for a set of frequencies, with a fixed GW source of SNR 3 and sky coordinates (RA, Dec) = (6.2 h, 45°). The error bars represent the Poisson fluctuations around line curves. As expected, FAR decreases with an increasing threshold, while FDR increases. FAR and FDR are computed from 10^4 noise simulations, with the fixed GW source added, based on the definitions in Eq. (14).

Next, we present the frequency-direction-averaged FAR and FDR for a range of injection SNRs in the third column of Fig. 5. We evaluate this for seven different injection SNRs (ranging from 1 to 4 in steps of 0.5). The average is performed for four different directions (RA, Dec) = (3 h, 87°), (0 h, 2.4°), (6.2 h, 30°), and (18.2 h, -30°) and seven frequencies (30, 60, 100, 200, 400, 700, and 1000 Hz). The error bars represent fluctuations around the solid curve across different frequencies and sky directions. The Peak SNR threshold decreases from left to right along the curve. Notably, FDR is dependent on source power and decreases as the injection SNR (or source power) increases.

In hypothesis testing, several criteria exist for determining an “optimal” threshold on the test statistics to accept or reject hypotheses. The Neyman-Pearson criterion, as outlined in the literature, suggests setting the threshold on the test statistics such that the FDR is minimized while keeping the FAR fixed [17]. In GWB analyzes, the FAR is typically set to 5%. The plot in the middle panel of Fig. 5 will be helpful for determining an appropriate threshold for the Peak SNR statistic.

In the next section, we will demonstrate the improve-

ments that the Peak SNR statistic brings to the identification of sub-threshold follow-up candidates.

3. Impact on FDR with Follow-up of Sub-threshold Candidates

The sensitivity of the ASAF radiometer search can be enhanced by following up on sub-threshold candidates using searches optimized for continuous wave signals. In the O3 ASAF search [30], 515 candidates were identified [30]. However, no significant candidates were found after a follow-up using the Viterbi algorithm [35].

In this section, we compare the Peak Finder method with the Full-Sky method. The FDR here is defined differently from Eq. (14) as ROC curves are not informative if we want to compare the two methods due to two reasons. First, the definition in Eq. (14) does not account for the difference in the number of trial factors of these two methods as shown in Fig. 3. So, we cannot directly compare them. Secondly, it is very difficult to compare the ROC curves of the two methods since the points on the ROC curve with same SNR threshold do not line up together, either along the FAR axis or the FDR axis. Thus, we replace the SNR-based FDR definition with one based on the number of candidates that can be followed up, N (Eq. (15)).

In this section, we simulate $N_{\text{sim}} = 10^4$ noise realizations and add a fixed GW source with parameters chosen from the combinations mentioned in previous section. For each realization, we rank all pixels in the Full-sky method, or only the peaks in the Peak-Finder method, according to their SNR, and then select the top N pixels with the highest SNR. If any of those top candidates lies within the diffraction-limited region centered at the injection direction, we consider it a successful detection. Given N_{det} is the number of realizations with a successful detection, we define the FDR as follows:




FIG. 6: A comparison of the False Dismissal Rate between the Full-Sky (dashed) and Peak-Finder (solid) methods as a function of the number of follow-up candidates. In each case, FDR decreases significantly when $N > 1$ with the Peak-Finder method. For example, in the left panel at 30 Hz, FDR is reduced by approximately a factor of 3 for $N = 2$. As the frequency increases, the difference in FDR between the two methods diminishes, and beyond a certain point, the trend reverses—FDR for the Full-Sky method becomes lower than that of the Peak-Finder method. Moving from the left to the right panel, the injection SNR increases, while the top and bottom rows correspond to injections in two distinct sky directions. As the injection SNR increases, the overall FDR decreases while maintaining the same trend with frequency variations.

$$\text{FDR} = 1 - \frac{N_{\text{det}}}{N_{\text{sim}}}. \quad (15)$$

In Fig. 6, we present the variation of FDR with the number of candidates for a set of frequencies (30, 100, 300, 400, 700, 1000 Hz), injection SNRs (3, 4, 5), and sky directions (RA, Dec) = (6.2h, 45°) and (18.2 h, -45°). The solid curve represents the FDR with Peak SNR statistic. The dashed curve corresponds to the FDR computed using the Full-Sky method.

Let us consider the top-left panel of Fig. 6 at $f = 30$ Hz. If we fix $N = 1$, then the FDR is the same for both methods. However, as N increases, the FDR in the Peak Finder method decreases significantly. For example, at $N = 2$, the FDR reduces from 0.19 (Full-Sky) to 0.04 (Peak Finder) by a factor of almost 5. However, as the frequency increases, the difference in FDR between the two methods starts to decrease up to 200 Hz. Beyond this, for frequencies of 400 and 700 Hz, the FDR obtained using the Peak SNR statistic becomes higher than that of the Full-Sky method. The bottom-left panel shows similar plots for a different direction (RA, Dec) = (21 h, -45°). As the injection SNR increases from left to right, the FDR behavior remains the same as for SNR = 3, but overall, the FDR decreases as the SNR increases.

We note that, at each frequency, the FDR for the Peak SNR statistic initially decreases sharply, and then flattens out. In contrast, the FDR for the Full-Sky statistic decreases more gradually, but continues to decline steadily. As a result, the Peak SNR FDR is lower than the Full-Sky FDR, for small N and trend reverses at higher N . This is particularly evident for an injection SNR of 5. However, the trend on high N is not a concern, since, in practice, only a small number of candidates per frequency can be followed up making the effect at large N irrelevant.

Now we can utilize the Peak SNR statistic, which is immune to correlation among sky directions, to identify the top, say $x\%$ or N , candidates. The reduction in trial factors and the computational savings from identifying candidates linked to the same root cause may allow us to follow up on more uncorrelated samples, ultimately leading to a significant decrease in the FDR.

If we decide to follow only one candidate per frequency, the number of follow-up candidates identified could be increased using the method described in the [26, 30] search paper. Alternatively, if we allow any number of candidates from any frequency, then $x\%$ from all frequency-peak pairs, which are immune to correlated samples, can be used to identify candidates following the method prescribed in the previous section.

The improvement in the detection rate in realistic settings, along with its effect on strain sensitivity, will be explored in future work.

V. CONCLUSION AND FUTURE PROSPECTS

With the initiation of the first all-sky all-frequency radiometer search, a significant advancement has been made toward a fast and model-independent search for narrowband and persistent gravitational wave sources. However, several open questions remain in order to fully utilize the search's potential. In this work, we addressed the issue of correlated sky samples, which can impact candidate selection and follow-up efficiency.

In this study, we proposed a method to bundle correlated samples together and identify their representative using the Peak Finder algorithm. We then characterized the statistical properties of the Peak SNR statistic through Monte Carlo simulations.

We analyzed the FDR and FAR as a function of the detection threshold on the Peak SNR statistic in ASAF radiometer search. We demonstrated that the Peak SNR method, which is immune to correlations among sky directions, can effectively reduce trial factors and computational costs while improving candidate selection for follow-up analysis. Our results show that at lower frequencies (e.g., 30 Hz), the Peak SNR method significantly reduces the FDR when multiple candidates per frequency bin ($N > 1$) are considered. However, as frequency increases beyond 200 Hz, the difference between the Peak SNR and Full-Sky methods diminishes. This suggests that the optimal follow-up strategy may depend on frequency and signal characteristics. Future work will explore the impact of these findings in realistic search settings, including improvements in detection rates and strain sensitivity.

Here, we have proposed a simple Peak Finder method to mitigate the issue of correlated sky samples in ASAF radiometer searches. In the future, more sophisticated clustering methods could be developed or applied to further refine candidate selection and improve detection efficiency.

ACKNOWLEDGMENTS

We would like to thank Iuri La Rosa for carefully reviewing this manuscript as part of LVK internal review process. This research has made use of folded data obtained from the Zenodo [53], provided by the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA. This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation. The authors are grateful for computational resources provided by the Inter-University Center for Astronomy and Astrophysics (IUCAA), Pune, India and LIGO Laboratory and

supported by National Science Foundation Grants PHY-0757058 and PHY-0823459. This article has a LIGO document number LIGO-P2500691.

A. S. acknowledges the use of the Maple cluster at the University of Mississippi (funded by NSF Grant CHE-1338056). D. A. acknowledges financial support from the IUCAA, Pune, India, during the initial phase and the Actions de Recherche Concertées (ARC) and Le Fonds spécial pour la recherche (FSR) of the Fédération Wallonie-Bruxelles during the intermediate phase and The University of Texas Rio Grande Valley, Texas, USA during final phase of this work. S. M. acknowledges the Department of Science and Technology (DST), Ministry of Science and Technology, India, for the support provided under the esteemed Swarna Jayanti Fellowships scheme. We used numerous software packages such as NumPy [54], SciPy [55], healpy and HEALPix [33, 34], h5py [56], PyStoch [32, 57], pandas [58, 59], and MATPLOTLIB [60] in this work.

Appendix A: Beam Matrix/Point Spread Function and Noise Covariance Matrix

1. Mean and Covariance of SNR

We suppress the f dependence. We note that SNR is defined as

$$\hat{\rho}_\alpha \equiv \frac{\hat{D}_\alpha}{\sigma_\alpha}, \quad (\text{A1})$$

where $\hat{D}_\alpha = [\Gamma_{\alpha\alpha}]^{-1} X_\alpha$, $\sigma_\alpha = [\Gamma_{\alpha\alpha}]^{-1/2}$.

Let us understand the contribution of GW source and noise to dirty map, i.e., given [20]

$$\begin{aligned} X_\alpha(f) &= \tau \Delta f \operatorname{Re} \sum_{It} \frac{\gamma_\alpha^{I*}(t, f) C^I(t, f)}{P_{I_1}(t, f) P_{I_2}(t, f)}, \\ \Gamma \equiv \Gamma_{\alpha\beta}(f) &= \frac{\tau \Delta f}{2} \operatorname{Re} \sum_{It} \frac{\gamma_\alpha^{I*}(t, f) \gamma_\beta^I(t, f)}{P_{I_1}(t, f) P_{I_2}(t, f)}, \\ \langle X_\alpha(f) \rangle_{h,n} &= \tau \Delta f \operatorname{Re} \sum_{It} \frac{\gamma_\alpha^{I*}(t, f) \langle C^I(t, f) \rangle_{h,n}}{P_{I_1}(t, f) P_{I_2}(t, f)}, \\ &\propto \tau \Delta f \operatorname{Re} \sum_{It} \frac{\gamma_\alpha^{I*}(t, f)}{P_{I_1}(t, f) P_{I_2}(t, f)}, \\ \langle [\tilde{h}_{I_1}^* \tilde{h}_{I_2} + \tilde{n}_{I_1}^* \tilde{h}_{I_2} + \tilde{h}_{I_1}^* \tilde{n}_{I_2} + \tilde{n}_{I_1}^* \tilde{n}_{I_2}] \rangle_{h,n} & \\ &= \tau \Delta f \operatorname{Re} \sum_{It} \frac{\gamma_\alpha^{I*}(t, f)}{P_{I_1}(t, f) P_{I_2}(t, f)} \langle \tilde{h}_{I_1}^* \tilde{h}_{I_2} \rangle_{h,n} \\ &\equiv \sum_{\beta} \Gamma_{\alpha\beta}(f) \mathcal{P}_\beta(f). \end{aligned} \quad (\text{A2})$$

Then, the mean of SNR in presence of a point source in the direction α_0 , i.e., $\mathcal{P}_\alpha(f) = \delta_{\alpha\alpha_0} \mathcal{P}_{\alpha_0}(f)$ is given by,

$$\langle \hat{\rho}_\alpha(f) \rangle_{h,n} = \frac{\Gamma_{\alpha\alpha_0}(f) \mathcal{P}_{\alpha_0}(f)}{[\Gamma_{\alpha\alpha}]^{-1/2}}. \quad (\text{A3})$$

Owing to the finite width of $\Gamma_{\alpha\alpha_0}(f)$, the response to a point source is not only localized to pixel α_0 .

Next, when there is no GW signal in data then mean of SNR is zero (in presence of uncorrelated Gaussian noise). However, the pixel-to-pixel covariance of SNR (in weak signal and large chunk limit) is given by [20]

$$\begin{aligned} \langle \hat{\rho}_\alpha \hat{\rho}_\beta \rangle_{h,n} &\propto \frac{1}{[\Gamma_{\alpha\alpha}]^{-1/2} [\Gamma_{\beta\beta}]^{-1/2}} \\ &\sum_{II'tt'} \frac{\gamma_\alpha^{I*}(t, f) \gamma_\beta^I(t', f)}{P_{I_1}(t, f) P_{I_2}(t, f) P_{I'_1}(t', f) P_{I'_2}(t', f)} \\ &\langle \tilde{n}_{I_1}^*(t; f) \tilde{n}_{I_2}(t; f) \tilde{n}_{I'_1}(t'; f) \tilde{n}_{I'_2}^*(t'; f) \rangle_{h,n} \\ &= \frac{\Gamma_{\alpha\beta}}{[\Gamma_{\alpha\alpha}]^{-1/2} [\Gamma_{\beta\beta}]^{-1/2}}. \end{aligned} \quad (\text{A4})$$

We note that purely due to the map-making method (and finite width of detector beam response), the SNR estimators could have pixel-to-pixel correlation in the presence on only detector noise.

2. Diffraction limit

To get a quantitative estimate of pixel-to-pixel correlation, let us assume that detector noise is stationary

during a day and detector response is isotropic. Then,

$$\Gamma_{\alpha\alpha'} \propto \sum_t e^{i2\pi f(\hat{\Omega}_\alpha - \hat{\Omega}_{\alpha'}) \cdot \Delta \vec{x}_I(t)/c}. \quad (\text{A5})$$

Let us make another simplification: the projection of baseline length vector $\Delta \vec{x}_I(t)$ on $(\hat{\Omega}_\alpha - \hat{\Omega}_{\alpha'})$ vary from $-L/2$ to $L/2$ as earth rotates where L is the maximum distance between two detectors. The sum over time segments can be replaced by integration over length, then

$$\Gamma_{\alpha\alpha'} \propto \text{sinc} \left[\pi f (|\hat{\Omega}_\alpha - \hat{\Omega}_{\alpha'}|) \cdot L/c \right], \quad (\text{A6})$$

and first zero of “sinc” function sets condition on

$$|\hat{\Omega}_\alpha - \hat{\Omega}_{\alpha'}| = \frac{c}{fL}.$$

Appendix B: Algorithm Runtime

The Peak Finder algorithm takes between 3-5 milliseconds to find the peaks in a map on a laptop², as shown in Fig. 7. Here we have plotted the mean runtime and its standard deviation for over 10^4 maps at each frequency.




FIG. 7: Time taken for the peak finder algorithm to identify peaks as a function of frequency.

-
- [1] A. G. Abac et al. GWTC-4.0: Updating the Gravitational-Wave Transient Catalog with Observations from the First Part of the Fourth LIGO-Virgo-KAGRA Observing Run. 8 2025.
 - [2] Magdalena Sieniawska and Michał Bejger. Continuous gravitational waves from neutron stars: Current status and prospects. *Universe*, 5(11), 2019. ISSN 2218-1997. doi:10.3390/universe5110217. URL <https://www.mdpi.com/2218-1997/5/11/217>.
 - [3] Rodrigo Tenorio, David Keitel, and Alicia M. Sintes. Search methods for continuous gravitational-wave signals from unknown sources in the advanced-detector era. *Universe*, 7(12), 2021. ISSN 2218-1997. doi: 10.3390/universe7120474. URL <https://www.mdpi.com/2218-1997/7/12/474>.
 - [4] Ornella Juliana Piccinni. Status and Perspectives of Continuous Gravitational Wave Searches. *Galaxies*, 10(3):72, 2022. doi:10.3390/galaxies10030072.
 - [5] Keith Riles. Searches for continuous-wave gravitational radiation. *Living Rev. Rel.*, 26(1):3, 2023. doi: 10.1007/s41114-023-00044-3.
 - [6] Karl Wette. Searches for continuous gravitational waves from neutron stars: A twenty-year retrospective. *Astropart. Phys.*, 153:102880, 2023. doi: 10.1016/j.astropartphys.2023.102880.

² Laptop specifications: Apple’s ARM-based M1 Pro chip, with 10 CPU cores (8 performance cores clocked at 3.2 GHz and 2 efficiency cores at 2 GHz) and 16 GB of memory.

- [7] B. Haskell and M. Bejger. Astrophysics with continuous gravitational waves. *Nature Astron.*, 7(10):1160–1170, 2023. doi:10.1038/s41550-023-02059-w.
- [8] Bruce Allen. The Stochastic gravity wave background: Sources and detection. In *Les Houches School of Physics: Astrophysical Sources of Gravitational Radiation*, pages 373–417, 4 1996.

- [9] Nelson Christensen. Stochastic Gravitational Wave Backgrounds. *Rept. Prog. Phys.*, 82(1):016903, 2019. doi: 10.1088/1361-6633/aae6b5.
- [10] Tania Regimbau. The quest for the astrophysical gravitational-wave background with terrestrial detectors. *Symmetry*, 14(2), 2022. ISSN 2073-8994. doi: 10.3390/sym14020270. URL <https://www.mdpi.com/2073-8994/14/2/270>.
- [11] Chiara Caprini and Daniel G. Figueroa. Cosmological Backgrounds of Gravitational Waves. *Class. Quant. Grav.*, 35(16):163001, 2018. doi:10.1088/1361-6382/aac608.
- [12] Joseph D. Romano and Neil J. Cornish. Detection methods for stochastic gravitational-wave backgrounds: a unified treatment. *Living Reviews in Relativity*, 20(1), apr 2017. doi:10.1007/s41114-017-0004-1. URL <https://doi.org/10.1007%2Fs41114-017-0004-1>.
- [13] R. Abbott et al. Diving below the spin-down limit: Constraints on gravitational waves from the energetic young pulsar psr j0537-6910. *The Astrophysical Journal Letters*, 913(2):L27, may 2021. doi: 10.3847/2041-8213/abffcd. URL <https://dx.doi.org/10.3847/2041-8213/abffcd>.
- [14] R. Abbott et al. Searches for gravitational waves from known pulsars at two harmonics in the second and third ligo-virgo observing runs. *The Astrophysical Journal*, 935(1):1, aug 2022. doi:10.3847/1538-4357/ac6acf. URL <https://dx.doi.org/10.3847/1538-4357/ac6acf>.
- [15] Nelson Christensen. Stochastic gravitational wave backgrounds. *Reports on Progress in Physics*, 82(1):016903, nov 2018. doi:10.1088/1361-6633/aae6b5. URL <https://dx.doi.org/10.1088/1361-6633/aae6b5>.
- [16] Neil J. Cornish and Joseph D. Romano. When is a gravitational-wave signal stochastic? *Phys. Rev. D*, 92:042001, Aug 2015. doi:10.1103/PhysRevD.92.042001. URL <https://link.aps.org/doi/10.1103/PhysRevD.92.042001>.
- [17] Bruce Allen and Joseph D. Romano. Detecting a stochastic background of gravitational radiation: Signal processing strategies and sensitivities. *Phys. Rev. D*, 59:102001, Mar 1999. doi: 10.1103/PhysRevD.59.102001. URL <https://link.aps.org/doi/10.1103/PhysRevD.59.102001>.
- [18] Bruce Allen and Adrian C. Ottewill. Detection of anisotropies in the gravitational wave stochastic background. *Phys. Rev. D*, 56:545–563, 1997. doi: 10.1103/PhysRevD.56.545.
- [19] Stefan W Ballmer. A radiometer for stochastic gravitational waves. *Classical and Quantum Gravity*, 23(8): S179, mar 2006. doi:10.1088/0264-9381/23/8/S23. URL <https://dx.doi.org/10.1088/0264-9381/23/8/S23>.
- [20] Sanjit Mitra, Sanjeev Dhurandhar, Tarun Souradeep, Albert Lazzarini, Vuk Mandic, Sukanta Bose, and Stefan Ballmer. Gravitational wave radiometry: Mapping a stochastic gravitational wave background. *Physical Review D*, 77(4), feb 2008. doi:10.1103/physrevd.77.042002. URL <https://doi.org/10.1103%2Fphysrevd.77.042002>.
- [21] B. Abbott, R. Abbott, et al. Upper limit map of a background of gravitational waves. *Phys. Rev. D*, 76: 082003, Oct 2007. doi:10.1103/PhysRevD.76.082003. URL <https://link.aps.org/doi/10.1103/PhysRevD.76.082003>.
- [22] J. Abadie, B. P. Abbott, et al. Directional limits on persistent gravitational waves using ligo s5 science data. *Phys. Rev. Lett.*, 107:271102, Dec 2011. doi: 10.1103/PhysRevLett.107.271102. URL <https://link.aps.org/doi/10.1103/PhysRevLett.107.271102>.
- [23] B. P. Abbott, R. Abbott, et al. Directional limits on persistent gravitational waves from advanced ligo’s first observing run. *Phys. Rev. Lett.*, 118:121102, Mar 2017. doi: 10.1103/PhysRevLett.118.121102. URL <https://link.aps.org/doi/10.1103/PhysRevLett.118.121102>.
- [24] B. P. Abbott and Others. Directional limits on persistent gravitational waves using data from Advanced LIGO’s first two observing runs. *Phys. Rev. D*, 100:062001, Sep 2019. doi:10.1103/PhysRevD.100.062001. URL <https://link.aps.org/doi/10.1103/PhysRevD.100.062001>.
- [25] R. Abbott and Others. Search for anisotropic gravitational-wave backgrounds using data from Advanced LIGO and Advanced Virgo’s first three observing runs. *Phys. Rev. D*, 104:022005, Jul 2021. doi: 10.1103/PhysRevD.104.022005. URL <https://link.aps.org/doi/10.1103/PhysRevD.104.022005>.
- [26] A. G. Abac et al. Directional Search for Persistent Gravitational Waves: Results from the First Part of LIGO-Virgo-KAGRA’s Fourth Observing Run. 10 2025.
- [27] C. Messenger, H. J. Bulten, S. G. Crowder, V. Dergachev, D. K. Galloway, E. Goetz, R. J. G. Jonker, P. D. Lasky, G. D. Meadors, A. Melatos, S. Premachandra, K. Riles, L. Sammut, E. H. Thrane, J. T. Whelan, and Y. Zhang. Gravitational waves from scorpius x-1: A comparison of search methods and prospects for detection with advanced detectors. *Phys. Rev. D*, 92:023006, Jul 2015. doi:10.1103/PhysRevD.92.023006. URL <https://link.aps.org/doi/10.1103/PhysRevD.92.023006>.
- [28] Patrick Meyers. *Cross-correlation searches for persistent gravitational waves with Advanced LIGO and noise studies for current and future ground-based gravitational-wave detectors*. PhD thesis, University of Minnesota, 2018. URL <https://conservancy.umn.edu/handle/11299/199038>.
- [29] Claudio Salvatore, Iuri La Rosa, Paola Leaci, Francesco Amicucci, Pia Astone, Sabrina D’Antonio, Luca D’Onofrio, Cristiano Palomba, Lorenzo Pierini, and Francesco Safai Tehrani. Harnessing the potential of pys-toch: Detecting continuous gravitational waves from interesting supernova remnant targets. *Phys. Rev. D*, 112(8):083051, 2025. doi:10.1103/xydb-k4vw.
- [30] R. Abbott et al. 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(12):122001, 2022. doi: 10.1103/PhysRevD.105.122001.
- [31] Anirban Ain, Prathamesh Dalvi, and Sanjit Mitra. Fast gravitational wave radiometry using data folding. *Physical Review D*, 92(2), jul 2015. doi: 10.1103/physrevd.92.022003. URL <https://doi.org/10.1103%2Fphysrevd.92.022003>.
- [32] Anirban Ain, Jishnu Suresh, and Sanjit Mitra. Very fast stochastic gravitational wave background map making using folded data. *Physical Review D*, 98(2), jul 2018. doi:10.1103/physrevd.98.024001. URL <https://doi.org/10.1103%2Fphysrevd.98.024001>.
- [33] 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. *The Astrophysical Journal*, 622(2):759, apr 2005. doi: 10.1086/427976. URL <https://dx.doi.org/10.1086/427976>.
- [34] Andrea Zonca, Leo P. Singer, Daniel Lenz, Martin Reinecke, Cyrille Rosset, Eric Hivon, and Krzysztof M. Gorski. healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in python. *Journal of Open Source Software*, 4(35):1298, 2019. doi: 10.21105/joss.01298. URL <https://doi.org/10.21105/joss.01298>.
- [35] Alan M. Knee, Helen Du, Evan Goetz, Jess McIver, Julian B. Carlin, Ling Sun, Liam Dunn, Lucy Strang, Hannah Middleton, and Andrew Melatos. Search for continuous gravitational waves directed at subthreshold radiometer candidates in o3 ligo data. *Phys. Rev. D*, 109:062008, Mar 2024. doi: 10.1103/PhysRevD.109.062008. URL <https://link.aps.org/doi/10.1103/PhysRevD.109.062008>.
- [36] Erik Floden, Vuk Mandic, Andrew Matas, and Leo Tsukada. Angular resolution of the search for anisotropic stochastic gravitational-wave background with terrestrial gravitational-wave detectors. *Phys. Rev. D*, 106: 023010, Jul 2022. doi:10.1103/PhysRevD.106.023010. URL <https://link.aps.org/doi/10.1103/PhysRevD.106.023010>.
- [37] J. Aasi and Others. Directed search for continuous gravitational waves from the galactic center. *Phys. Rev. D*, 88:102002, Nov 2013. doi: 10.1103/PhysRevD.88.102002. URL <https://link.aps.org/doi/10.1103/PhysRevD.88.102002>.
- [38] Berit Behnke, Maria Alessandra Papa, and Reinhard Prix. Postprocessing methods used in the search for continuous gravitational-wave signals from the galactic center. *Phys. Rev. D*, 91:064007, Mar 2015. doi: 10.1103/PhysRevD.91.064007. URL <https://link.aps.org/doi/10.1103/PhysRevD.91.064007>.
- [39] Avneet Singh, Maria Alessandra Papa, Heinz-Bernd Eggenstein, and Sinéad Walsh. Adaptive clustering procedure for continuous gravitational wave searches. *Phys. Rev. D*, 96:082003, Oct 2017. doi: 10.1103/PhysRevD.96.082003. URL <https://link.aps.org/doi/10.1103/PhysRevD.96.082003>.
- [40] B. Beheshtipour and M. A. Papa. Deep learning for clustering of continuous gravitational wave candidates. ii. identification of low-snr candidates. *Phys. Rev. D*, 103: 064027, Mar 2021. doi:10.1103/PhysRevD.103.064027. URL <https://link.aps.org/doi/10.1103/PhysRevD.103.064027>.
- [41] B. Steltner, T. Menne, M. A. Papa, and H.-B. Eggenstein. Density-clustering of continuous gravitational wave candidates from large surveys. *Phys. Rev. D*, 106: 104063, Nov 2022. doi:10.1103/PhysRevD.106.104063. URL <https://link.aps.org/doi/10.1103/PhysRevD.106.104063>.
- [42] Deepali Agarwal, Jishnu Suresh, Sanjit Mitra, and Anirban Ain. Angular power spectra of anisotropic stochastic gravitational wave background: Developing statistical methods and analyzing data from ground-based detectors. *Phys. Rev. D*, 108:023011, Jul 2023. doi: 10.1103/PhysRevD.108.023011. URL <https://link.aps.org/doi/10.1103/PhysRevD.108.023011>.
- [43] R. Abbott, T. D. Abbott, et al. Upper limits on the isotropic gravitational-wave background from advanced ligo and advanced virgo's third observing run. *Phys. Rev. D*, 104:022004, Jul 2021. doi: 10.1103/PhysRevD.104.022004. URL <https://link.aps.org/doi/10.1103/PhysRevD.104.022004>.
- [44] Eric Thrane, Stefan Ballmer, Joseph D. Romano, Sanjit Mitra, Dipongkar Talukder, Sukanta Bose, and Vuk Mandic. Probing the anisotropies of a stochastic gravitational-wave background using a network of ground-based laser interferometers. *Phys. Rev. D*, 80: 122002, Dec 2009. doi:10.1103/PhysRevD.80.122002. URL <https://link.aps.org/doi/10.1103/PhysRevD.80.122002>.
- [45] A. Lazzarini and J. Romano. Use of overlapping windows in the stochastic background search. *Technical Report No. LIGO-T040089-x0*, 2004. URL <https://dcc.ligo.org/LIGO-T040089/public>.
- [46] Sambit Panda, Swetha Bhagwat, Jishnu Suresh, and Sanjit Mitra. Stochastic gravitational wave background mapmaking using regularized deconvolution. *Phys. Rev. D*, 100:043541, Aug 2019. doi: 10.1103/PhysRevD.100.043541. URL <https://link.aps.org/doi/10.1103/PhysRevD.100.043541>.
- [47] Deepali Agarwal, Jishnu Suresh, Sanjit Mitra, and Anirban Ain. Upper limits on persistent gravitational waves using folded data and the full covariance matrix from advanced LIGO's first two observing runs. *Physical Review D*, 104(12), dec 2021. doi: 10.1103/physrevd.104.123018. URL <https://doi.org/10.1103/PhysRevD.104.123018>.
- [48] Liting Xiao, Arianna I. Renzini, and Alan J. Weinstein. Model-independent search for anisotropies in stochastic gravitational-wave backgrounds and application to ligo-virgo's first three observing runs. *Phys. Rev. D*, 107: 122002, Jun 2023. doi:10.1103/PhysRevD.107.122002. URL <https://link.aps.org/doi/10.1103/PhysRevD.107.122002>.
- [49] Eric Thrane, Sanjit Mitra, Nelson Christensen, Vuk Mandic, and Anirban Ain. All-sky, narrowband, gravitational-wave radiometry with folded data. *Phys. Rev. D*, 91:124012, Jun 2015. doi:10.1103/PhysRevD.91.124012. URL <https://link.aps.org/doi/10.1103/PhysRevD.91.124012>.
- [50] Sanjeev Dhurandhar, Badri Krishnan, Himan Mukhopadhyay, and John T. Whelan. Cross-correlation search for periodic gravitational waves. *Phys. Rev. D*, 77:082001, Apr 2008. doi:10.1103/PhysRevD.77.082001. URL <https://link.aps.org/doi/10.1103/PhysRevD.77.082001>.
- [51] John T. Whelan, Santosh Sundaresan, Yuanhao Zhang, and Prabath Peiris. Model-based cross-correlation search for gravitational waves from scorpius x-1. *Phys. Rev. D*, 91:102005, May 2015. doi: 10.1103/PhysRevD.91.102005. URL <https://link.aps.org/doi/10.1103/PhysRevD.91.102005>.
- [52] E Goetz and K Riles. An all-sky search algorithm for continuous gravitational waves from spinning neutron stars in binary systems. *Classical and Quantum Gravity*, 28(21):215006, sep 2011. doi:10.1088/0264-9381/28/21/215006. URL <https://dx.doi.org/10.1088/0264-9381/28/21/215006>.
- [53] Folded data for first three observing runs of Advanced LIGO and Advanced Virgo, March 2022. URL <https://zenodo.org/records/6326656>.

- [54] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Shepard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, September 2020. doi:10.1038/s41586-020-2649-2. URL <https://doi.org/10.1038/s41586-020-2649-2>.
- [55] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi:10.1038/s41592-019-0686-2.
- [56] Tim Jenness. Reimplementing the hierarchical data system using hdf5. *Astronomy and Computing*, 12: 221–228, September 2015. ISSN 2213-1337. doi: 10.1016/j.ascom.2015.02.003. URL <http://dx.doi.org/10.1016/j.ascom.2015.02.003>.
- [57] Anirban Ain, Jishnu Suresh, Sudhagar Suyamprakasam, and Sanjit Mitra. PyStoch Code. https://git.ligo.org/stochastic-public/stochastic/-/tree/master/PyStoch?ref_type=heads, 2021.
- [58] The pandas development team. pandas-dev/pandas: Pandas, February 2020. URL <https://doi.org/10.5281/zenodo.3509134>.
- [59] Wes McKinney. Data Structures for Statistical Computing in Python. In Stéfan van der Walt and Jarrod Millman, editors, *Proceedings of the 9th Python in Science Conference*, pages 56 – 61, 2010. doi:10.25080/Majora-92bf1922-00a.
- [60] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. doi:10.1109/MCSE.2007.55.