

Frequency-Resolved Atlas of the Sky in Continuous Gravitational Waves

Vladimir Dergachev^{1,2,*} and Maria Alessandra Papa^{1,2,3,†}

¹Max Planck Institute for Gravitational Physics (Albert Einstein Institute),
Callinstrasse 38, 30167 Hannover, Germany

²Leibniz Universität Hannover, D-30167 Hannover, Germany

³University of Wisconsin Milwaukee, 3135 North Maryland Avenue, Milwaukee, Wisconsin 53211, USA

 (Received 21 February 2022; revised 10 March 2023; accepted 10 April 2023; published 5 May 2023)

We present the first atlas of the continuous gravitational wave sky, produced using LIGO O3a public data. For each 0.045 Hz frequency band and every point on the sky the atlas provides gravitational wave amplitude upper limits, signal-to-noise ratios (SNRs), and frequencies where the search measures the maximum SNR. The approximately top 1.5% of the atlas results are reanalyzed with a series of more sensitive searches with the purpose of finding high SNR long coherence signals from isolated neutron stars. However, these searches do not reveal the presence of such signals. The results presented in the atlas are produced with the Falcon pipeline and cover nearly monochromatic gravitational-wave signals in the 500–1000 Hz band, with up to $\pm 5 \times 10^{-11}$ Hz/s frequency derivative. The Falcon pipeline computes loosely coherent power estimates to search data using a succession of coherence lengths. For this search we use six months of data, started with a 12 hour coherence length and progress to six days. Compared to the most sensitive results previously published (also produced with the Falcon pipeline), our upper limits are 50% more constraining. Neutron stars with ellipticity of 10^{-8} can be detected up to 150 pc away, while allowing for a large fraction of the stars' energy to be lost through nongravitational channels. These results are within an order of magnitude of the minimum neutron star ellipticity of 10^{-9} suggested by Woan *et al.* [*Astrophys. J. Lett.* **863**, L40 (2018)]

DOI: 10.1103/PhysRevX.13.021020

Subject Areas: Astrophysics, Gravitation

I. INTRODUCTION

Astronomy began with the observation of light from persistent sources. Gravitational-wave astronomy, on the contrary, started with the observation of signals from the merger of compact stars, which are loud transient events lasting less than a minute [1].

This paper is about a different type of gravitational-wave signals: those consistently “on” over months and even years. Such signals, commonly referred to as continuous gravitational waves, are expected from a variety of astrophysical scenarios and could be detected by the LIGO, Virgo, or KAGRA interferometers [2–4]. Among the scenarios that could give rise to continuous waves detectable by ground-based detectors, there are rotating neutron stars emitting due to a sustained quadrupolar deformation,

fluid oscillations [5–7], as well as more exotic scenarios as superradiance from ultralight bosons around spinning black holes [8–10].

The detection of a continuous gravitational wave will lead to high-precision tests of general relativity and, depending on the emission mechanism, may probe the interior and physics of neutron stars in an entirely new way or unveil new physics [11]. For this reason there exist a variety of research efforts aiming at detecting such signals [12].

Computational searches combining data collected over weeks or even longer periods are carried out using a broad variety of approaches. This variety arises because the most sensitive methods cannot be used due to computational limitations, so different speed-optimized methods have been designed that sacrifice either breadth of search or sensitivity, or a bit of both. Our loosely coherent algorithms [13–15] and, in particular, the Falcon search pipeline [16–18] that is used for the analysis presented here, have performed very well in this trade-off.

In this paper we provide a window on the universe of continuous gravitational-wave sources. For the first time ever, every sky location has an associated spectrum of upper limits on the gravitational-wave amplitude, as well as a signal-to-noise ratio (SNR) spectrum. These are, nearly raw, minimally processed results and access to them will enable new and independent searches, as we explain in

*vladimir.dergachev@aei.mpg.de

†maria.alessandra.papa@aei.mpg.de

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by the Max Planck Society.

Sec. V. Together with the new Einstein@Home results [19], these are the most sensitive results to date.

As it is common in broad surveys for continuous gravitational waves, we follow up and investigate less than 2% of the results of our search—and such results are marked in the atlas. None of the follow-ups result in a conclusive association with astrophysical sources. The rest of the atlas remains unexplored. Other scientists can use the atlas in combination with electromagnetic and particle data [20–24], and catalogs [25,26], in search of significant associations. Many other studies are possible. In the Appendix, we provide an example of an interesting search that can be easily performed using only the atlas data.

The atlas contains nearly two billion records. To make it easily accessible, the atlas has been designed so that it can be analyzed on a small personal computer using the Mapped Vector Library (MVL) library [27,28]. We provide open-source software (compatible with Linux, Windows, and Mac OS) that produces sky maps for wide frequency bands, allows queries of the atlas by sky position and/or frequency, and performs computations using the full scan of the data. Effectively this enables searches within the covered parameter space with the full sensitivity of the Falcon pipeline, but without the need for a large computing cluster. Our hope is to make gravitational-wave studies and discoveries accessible to scientists with a modest computing budget.

In the following sections we describe the gravitational wave signals, the analysis method, the search setup (including the follow-up), and the results of the analysis.

II. CONTINUOUS GRAVITATIONAL-WAVE SOURCES

The ground-based gravitational-wave detectors presently in operation [2–4] are sensitive in the audio-frequency range up to a few kilohertz, with the lower bound at a few tens of hertz given by the steep “seismic wall.”

The seismic wall at low frequencies precludes the observation of continuous gravitational waves from the decay of the orbit of white dwarf binaries with large separation. For this reason the classical sources of audio-frequency continuous gravitational-wave signals are rotating neutron stars with nonaxisymmetric deformations. For such sources the gravitational-wave emission is powered by the stars’ rotation.

Theory predicts that the neutron star crust can support quadrupolar deformations of 10^{-6} and perhaps even larger for neutron stars with exotic composition; however, signals like this have not yet been detected [29–31].

A constraint on the likelihood of the existence of rapidly spinning neutron stars with deformations much larger than 10^{-8} comes from the lifetime of sources: due to emission of the gravitational waves, the spin rate of the neutron star decreases in time. Since the energy loss is proportional to the sixth power of the rotation rate, highly deformed (highly emitting) high frequency sources are less likely to exist in our Galaxy because they quickly become

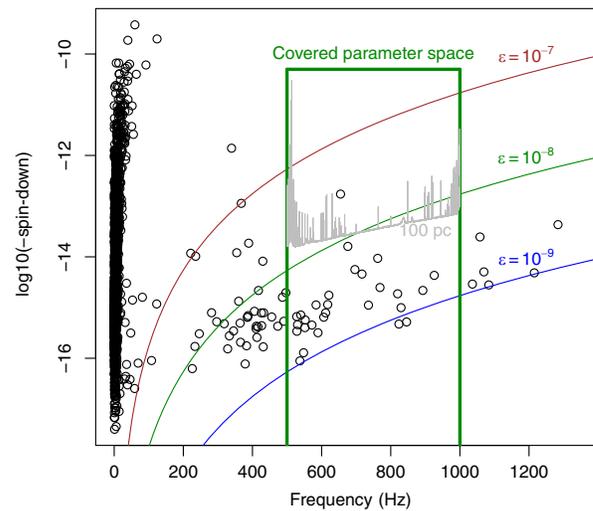


FIG. 1. This plot shows the frequency and frequency derivative of isolated neutron stars from the Australia Telescope National Facility Pulsar Catalogue [25] (circles). The pulsar frequency is shown multiplied by a factor of 2, to correspond to the expected gravitational-wave frequency. The curves show the frequency and frequency-derivative combinations for a gravitar (a neutron star losing energy solely due to gravitational-wave emission) for different values of the ellipticity. Following Ref. [33] we observe that known pulsars have spin-down values that lie at or above the curve $\epsilon = 10^{-9}$. For gravitars this would imply that the minimum ellipticity is around 10^{-9} . The region above the gray trace and enclosed by the box of green straight lines shows the frequency and frequency-derivative combinations of optimally oriented gravitars at 100 pc, that our search could detect. The frequency derivative of a detectable gravitar is proportional to the square of the distance, so if the distance increases by a factor of 10, correspondingly the gray trace moves up by a factor of 100.

slowly rotating neutron stars [32]. In fact, known radio pulsars rotating faster than 250 Hz (500 Hz gravitational-wave frequency) have frequency derivative less than 2×10^{-13} Hz/s [25] (Fig. 1), corresponding to maximum deformations of the order of 10^{-8} and below.

Examination of properties of known pulsars suggests that a population might exist whose spin evolution is governed by gravitational-wave emission, with typical ellipticities in the range of $\approx 10^{-9}$ – 10^{-8} [33] (see Fig. 1).

III. ANALYSIS METHOD

The expected signals are nearly monochromatic, with small frequency drifts in time. The frequency $f(t)$ of such signals is well described by Taylor polynomials with few components [17]. For our search the second order Taylor polynomial is sufficient:

$$f(t) = f_0 + (t - t_0)f_1 + (t - t_0)^2 f_2 / 2, \quad (1)$$

We call these “IT2” type signals, where “I” indicates an isolated source (no binary modulation), “T” indicates

Taylor polynomial model, and “2” gives the degree of the polynomial [17].

The signal at the detector is more complex as it presents frequency modulations due to the Doppler effect induced by the relative motion between the source and the receiver. For isolated neutron stars, which are the target of this paper, the largest Doppler shifts are due to Earth orbital motion [34,35], and are on the order of $10^{-4}f_0$, with f_0 being the signal frequency. Additionally, the signal also presents amplitude modulations due to the nonuniform beam-pattern response of the detector across the sky.

The first step of the analysis consists in heterodyning the data at a frequency close to the target signal frequencies and extracting a narrow band from the heterodyned data. This step is repeated over the different frequency subbands which compound the broad search band and it is efficiently accomplished in the Fourier domain.

After this transformation the data for the band of interest are a complex time series $\{z(t)\}$, down-sampled with respect to the original data. In this representation, the gravitational-wave signal has the form

$$S(t) = h_0 A(t) e^{i\Phi(t)}, \quad (2)$$

where h_0 is the intrinsic amplitude of the signal at the detector and the amplitude response $A(t)$ depends on the relative orientation of the source and the detector and hence changes as Earth rotates. The phase evolution term $\Phi(t)$ at the detector is due to the Doppler-shifted frequency evolution of Eq. (1): it is independent of orientation but varies with sky location, frequency, and frequency derivatives [34,35].

The input data are then the sum of the unknown signal and noise $n(t)$:

$$z(t) = h_0 A(t) e^{i\Phi(t)} + n(t). \quad (3)$$

The term $n(t)$ is a combination of fundamental noise sources and “technical” noise sources. The former include shot noise and radiation pressure on the instrument mirrors, and can be considered Gaussian. The technical noise sources are due to accidental couplings between noise sources external and internal to the detector and the gravitational wave output channels. These couplings are highly dependent on frequency and are nonstationary. Since a lot of effort has gone into reducing these couplings, there are many “clean” bands where these couplings are small enough to be practically negligible. Even in these clean bands, however, the $n(t)$ is in general not a stationary process—the reason being that the beam power in the detector arms varies in time affecting both the level of fundamental noise sources and the gain of the detector.

Next we consider data with the amplitude response and phase evolution terms taken out:

$$\tilde{z}(t) = \frac{z(t) e^{-i\Phi_0(t)}}{A(t)} = h_0 e^{i[\Phi(t) - \Phi_0(t)]} + \frac{n(t) e^{-i\Phi_0(t)}}{A(t)}. \quad (4)$$

$\Phi_0(t)$ is an assumed phase evolution, in general distinct from the true evolution $\Phi(t)$ when we search for unknown signals. We assume that the amplitude evolution for a signal with phase $\Phi_0(t)$ is practically equal to $A(t)$ of the signal with phase evolution $\Phi(t)$, which is correct for signals from sky locations within 1° of each other and the same gravitational-wave polarization. A generic sky position can also be considered, leading to the slightly different expression [14] actually used in the search, but here we make this simplifying assumption in order to illustrate the main idea.

We construct a linear time-domain filter \mathcal{L} , say, a bandpass filter, with wide enough bandwidth to leave invariant signals of the type

$$h_0 e^{i[\Phi(t) - \Phi_0(t)]} \quad (5)$$

and with as narrow bandwidth as desired to reject the maximum possible amount of noise.

The power $P = \|\mathcal{L}\tilde{z}\|^2$ can then be used to establish a limit on the strength h_0 of the gravitational-wave signal in the data. This is best achieved based on excess power measurements, i.e., by considering a set of power measurements for a set of phase evolution models Φ_0 and by setting an upper limit under the assumption that only a single signal is present [36].

In our approach there is an inherent trade-off between the maximum phase mismatch of phase evolutions $\{\Phi(t)\}$ from $\Phi_0(t)$ and the narrowness of the filter \mathcal{L} . The narrower the filter, the less noise it lets through and the more sensitive the search.

The structure of the manifold of phase evolutions $\{\Phi(t)\}$ used to design the filter is very important, but its description goes well beyond the scope of this paper [35].

An important characteristic of the manifold of phase evolutions is the maximum phase variation V over the timescale δ :

$$V(\delta) = \max_{|t_1 - t_2| < \delta} |\Phi(t_1) - \Phi_0(t_1) - \Phi(t_2) + \Phi_0(t_2)|. \quad (6)$$

When $V(\delta)$ is close to 0, the filter \mathcal{L} can be a simple averaging of the data. As the variation grows, the filter \mathcal{L} needs to include an explicit—and more computationally expensive—demodulation. As $V(\delta)$ exceeds π , the exponent $\exp i[\Phi(t) - \Phi_0(t)]$ changes sign. The value of δ when $V(\delta) = \pi$ is known as coherence length.

By picking \mathcal{L} as the projection onto a constant, we obtain a very narrow filter that admits only $\Phi(t) - \Phi_0(t) = \text{const}$ —this is known as a “fully coherent” search. At the other extreme, if \mathcal{L} is the identity operator which leaves its input unchanged, the power $P = \|\mathcal{L}\tilde{z}\|^2 = \|\tilde{z}\|^2$ discards the phase information and we arrive at a method with the shortest coherence length possible given the bandwidth of the heterodyne procedure used to obtain $z(t)$. The loosely coherent method [13–15] employed by Falcon uses the filter \mathcal{L} with a tunable coherence length [15].

In practice, in order to search for unknown sources anywhere on the sky, one needs to repeat the above procedure for many different assumed phase evolutions Φ_0 , corresponding to different source sky positions and signal frequency and frequency derivatives. We thus arrive at the fundamental difficulty of all-sky searches: the number of different phase evolutions that need to be considered, and hence the computational demands, scale with coherence length, which refines the resolution in sky, frequency, and frequency derivatives.

The Falcon pipeline used in this paper employs a loosely coherent algorithm [13–15] that finds the maximum power P for a set of assumed evolutions $\{\Phi_0^i(t)\}$. Operating on a set allows us to frame the problem as the optimization of a highly oscillatory function and take advantage of efficiencies of scale.

The typical analysis takes several months to complete on a facility of the scale of the ATLAS computing cluster [37], so with every new analysis we strive to better take advantage of existing data and computational infrastructure. While limits of computational hardware dictate the choice of the algorithm used, compilation of the code, scheduling, threading of the compute jobs on the physical cores, and even aspects such as data layout are carefully considered for every run because they can have a large impact on computational speed.

Based on the power P , the algorithm establishes upper limits [36] and SNRs for every analyzed portion of the sky and every frequency band. This dataset is quite large and in the past was considered impractical to record, so only summary information was presented, obtained by grouping the results by frequency.

For this search, instead, thanks to the use of Mapped Vector Library files, we could keep and distribute all the results—upper limits and SNRs, adding up to over 70 GB of data. The MVL technology was originally used to extract large amounts of engineering data describing the pipeline’s operation in real time. Only later it was realized that it is also perfectly suited for the scientific output of the pipeline. The key property of MVL files is that they are optimized for memory mapping. For example, with a notebook with 16 GB RAM and solid-state drive, one can open the atlas in its entirety and access any data by simple array subscription, as if the notebook had an order of magnitude more memory than it actually does.

IV. SEARCH

Our search uses data of the LIGO H1 and L1 interferometers, from the publicly released O3a set [38]. We limit our search to the LIGO detectors because no other detector comes within a factor of 2 of their sensitivity. We set O3b data [39] aside as an independent dataset for verification of any surviving outlier. This is a standard approach for searches over a very broad parameter space: due to the large trials’ factor, noise alone produces outliers that can

only be vetoed through verification on a different dataset (see, for example, Ref. [40]). We note that this procedure can be used because the same physical signal is assumed present in both datasets.

The O3a dataset is contaminated by loud short-duration excesses of noise [41–43]. To address this, we first prepare short Hann-windowed Fourier transforms (SFTs) with 0.25 s duration. We then monitor the frequency band 20–36 Hz for power excess that indicates noise contamination and exclude the affected SFTs from the analysis. The remaining SFTs are searched with the Falcon pipeline. The 20–36 Hz band is chosen because it is outside of the search range and heuristically we find it to be an effective “witness band” to identify spectral disturbances.

We search for nearly monochromatic signals of the IT2 type given by Eq. (1) with the signal frequency f_0 and its first derivative f_1 defined at GPS (Global Positioning System) epoch $t_0 = 1\,246\,070\,000$ (2019 Jul 2 02:33:02 UTC).

The search targets emission in the 500–1000 Hz frequency range, covering frequency derivatives $|f_1| \leq 5 \times 10^{-11}$ Hz/s and $|f_2| \leq 4 \times 10^{-20}$ Hz/s². Compared to our earlier analyses on O2 data [18], we use a more sensitive dataset and we search a spin-down range larger by a factor $\gtrsim 16$. This provides a larger margin ($\times 2.9$ at 1000 Hz) for Doppler shifts due to stellar motion and energy loss due to magnetic fields (Fig. 1).

The analyzed frequency band is chosen to make the first gravitational wave atlas simpler to use, while providing the most useful data. The lower cutoff of 500 Hz is picked to avoid numerous detector artifacts, while the upper cutoff of 1000 Hz is imposed to limit total data size, as the size of the atlas scales as cube of upper frequency.

The pipeline employs four stages (Table I) with coherence length increasing from 12 h to 6 days. Results exceeding the SNR threshold of one stage pass to the next stage. The upper limits and the atlas are constructed based on the results of the first two stages only. This simplifies the upper limit procedure without significant impact on the sensitivity. The last stage of the pipeline has long enough coherence time that signals that have survived up to that stage can now be identified using the data from the individual interferometers separately. This means that signal-consistency checks can be performed by comparing the results from each detector.

V. RESULTS

The Falcon pipeline computes upper limits on the intrinsic amplitude h_0 of continuous gravitational waveforms as a function of signal frequency and source sky position. We provide upper limits [44] for circularly polarized gravitational waves, which are generated by optimally oriented sources, as well as worst-case upper limits computed by maximizing over polarization. Such upper limits hold for sources with the most unfavorable orientation.

TABLE I. Parameters for each stage of the search. Stage 4 refines outlier parameters by using a denser sampling of parameter space, and then subjects them to an additional consistency check based on the analyses of individual interferometer data. The results from stages 1 and 2 are used to construct the atlas, stages 3 and 4 are solely outlier follow-ups.

Stage	Coherence length (days)	SNR threshold
1	0.5	6
2	1	8
3	2	8
4	6	12

The upper limits are part of the frequency-resolved atlas of the sky, which also includes the measured SNR as a function of signal frequency and source sky position. Specifically, the atlas contains sky-resolved data for every 45 mHz signal-frequency band. The sky resolution increases proportionally to frequency, and at 1000 Hz at any location in the sky there is a grid point within a 0.26° radius. For each sky pixel and frequency bin we also provide the highest SNR value and the corresponding frequency. Our simulations show that if the maximum SNR exceeds 9 at the closest sky grid point to the signal location, the signal frequency is within 0.6 mHz of the frequency of the maximum, with 95% confidence.

The upper limits are summarized in Figs. 2 and 3, by taking the maximum over the sky in 45 mHz frequency bands. This is the traditional way to present continuous wave upper limit results.

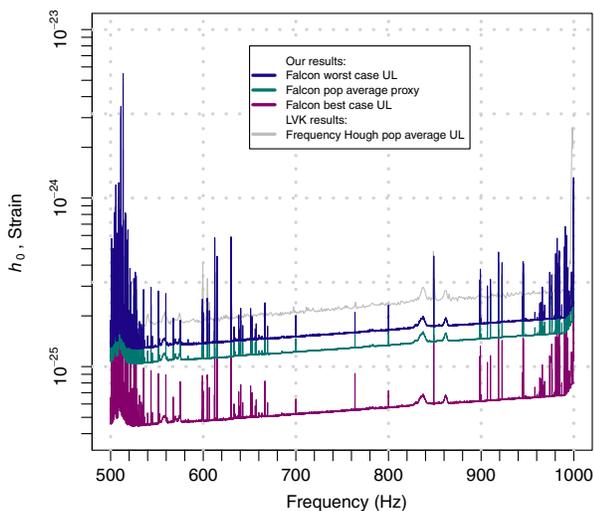


FIG. 2. Gravitational-wave intrinsic amplitude h_0 upper limits (UL) at 95% confidence as a function of signal frequency. The upper limits are a measure of the sensitivity of the search. We introduce a “population average” proxy upper limit in order to compare with the latest LIGO-Virgo all-sky results [45]. In this frequency range the results from Ref. [45] are a factor $\gtrsim 1.65$ less constraining than ours, albeit able to detect sources with much higher deformations.

The upper limits can be taken as a measure of the sensitivity of our search, and recast in terms of the equatorial ellipticity ε of sources that the search is able to detect:

$$\varepsilon = \frac{c^4}{4\pi^2 G} \frac{h_0 d}{I_{zz} f_0^2}, \quad (7)$$

where $I_{zz} = 10^{38} \text{ kg m}^2$ is the moment of inertia of the star with respect to the principal axis aligned with the rotation axis and d its distance.

This search is sensitive to sources with ellipticity of 10^{-8} up to 150 pc away, improving on our O2 results by a factor of ≈ 1.5 and on the full O3 results of Ref. [45] and Ref. [9] by a factor of ≈ 1.7 and 1.5, respectively. Signals sourced by ellipticities of 10^{-7} would have been detected from sources within a distance of 500 pc and, depending on orientation, as large as 1.5 kpc, at the highest frequencies. For lowest frequencies the distance is up to 200 pc, and as large as 600 pc for optimally oriented sources.

The results in Ref. [45] cover a much broader spin-down range than our search. They exclude neutron stars with ellipticities larger than 10^{-6} within a few kiloparsec, motivating a search like ours for lower ellipticity sources. We accomplish this with a spin-down range that covers ellipticities as high as 10^{-7} , while improving the sensitivity of the search.

The atlas is made possible due to our rigorous analytical upper limit procedure [36]. The analytical formula is required to efficiently produce billions of individual upper limits, and the statistical rigor is a necessity to ensure that any subset of

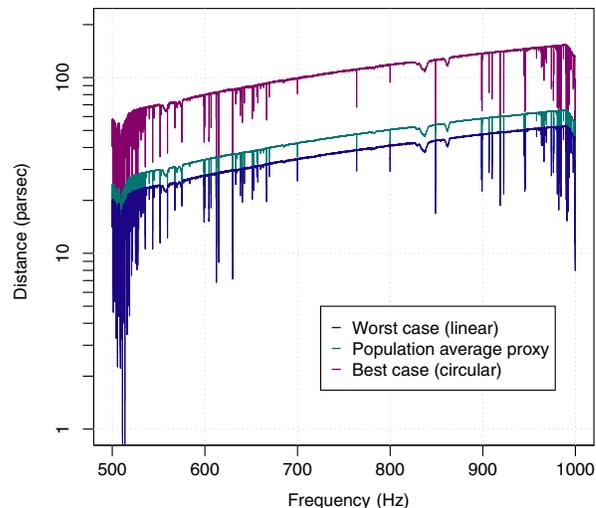


FIG. 3. Reach of the search for stars with ellipticity of 10^{-8} . The search is also sensitive to sources with ellipticities of 10^{-7} with a distance from Earth that is 10 times higher. The x axis is the gravitational wave frequency, which is twice the pulsar rotation frequency for emission due to an equatorial ellipticity. R modes and other emission mechanisms give rise to emission at different frequencies [7]. The top curve (purple) shows the reach for a population of circularly polarized sources; the middle curve (cyan) holds for a population of sources with random orientations; the bottom curve (blue) holds for linearly polarized sources.

results is individually valid regardless of any peculiarities in the distribution of underlying data. In contrast, many searches [9,12,19,45,46] use Monte Carlo simulations that establish upper limits over the whole sky. This method is impractical to use on many sky locations.

The most sensitive population-average all-sky upper limits in Ref. [45] did not use Monte Carlo simulations, but rather an analytical formula. While the authors provide a scaling prescription to derive upper limits at arbitrary sky positions, its validity is not obvious. In particular, it is not clear why the scaling factor derived from *average* antenna beam-pattern function over the population should be representative of the 95% *percentile* population amplitudes. This confusion may lead to an underestimate of the upper limits. Even for the all-sky case the formula is heuristic, at best. It was derived under the assumption of independent identically distributed Gaussian noise and neglects the variation of the noise level in real data. The deviations from Gaussianity are also ignored which precludes us from establishing rigorous upper limits.

A counterpart of our atlas is produced by the radiometer search for the gravitational-wave stochastic background [47]. Both the radiometer and this search are affected by loud continuous signals. However, this search is specifically tuned to continuous signals as it, for example, tracks the signal Doppler shifts due to the Earth motion over months of observation. This makes it significantly more sensitive (we estimate by at least a factor of 10) to continuous waves—and also more computationally expensive—than the radiometer search.

All-sky surveys for continuous gravitational waves are extremely computationally intensive: the results of this search cost in excess of 36×10^6 CPU hours. Our atlas [44] will allow others to benefit from this investment: astronomical observations of pulsations from a certain location in the sky may be checked against high SNR occurrences at the same location and consistent frequencies; the atlas can immediately be scanned across the whole frequency range at an interesting sky position, for example, that of a supernova remnant (as we demonstrate in the Appendix), in search of a significant result; with an estimate of the distance of a gravitational-wave source, upper limits on its ellipticity can be readily computed; models of neutron star populations can be directly convolved with the atlas upper limits to make predictions of their detectability, and the same holds for primordial black hole populations emitting continuous gravitational waves through their orbital energy loss [48], or for tens of solar mass black hole populations sourcing continuous waves through superradiance [8]. The fine granularity of our atlas (Figs. 4 and 5) produces high-number statistics and, in case of a coincidence with some other observation, this will greatly aid in the estimation of the chance probability of any finding.

Five outliers survive all stages (Table II). Three correspond to simulated signals [49] injected via radiation pressure on the detectors' test masses (Table III).

One is induced by a large instrumental artifact in H1 at 993.300 Hz. This instrumental artifact is quite loud, but is not included in the known line list [38]. The reanalysis

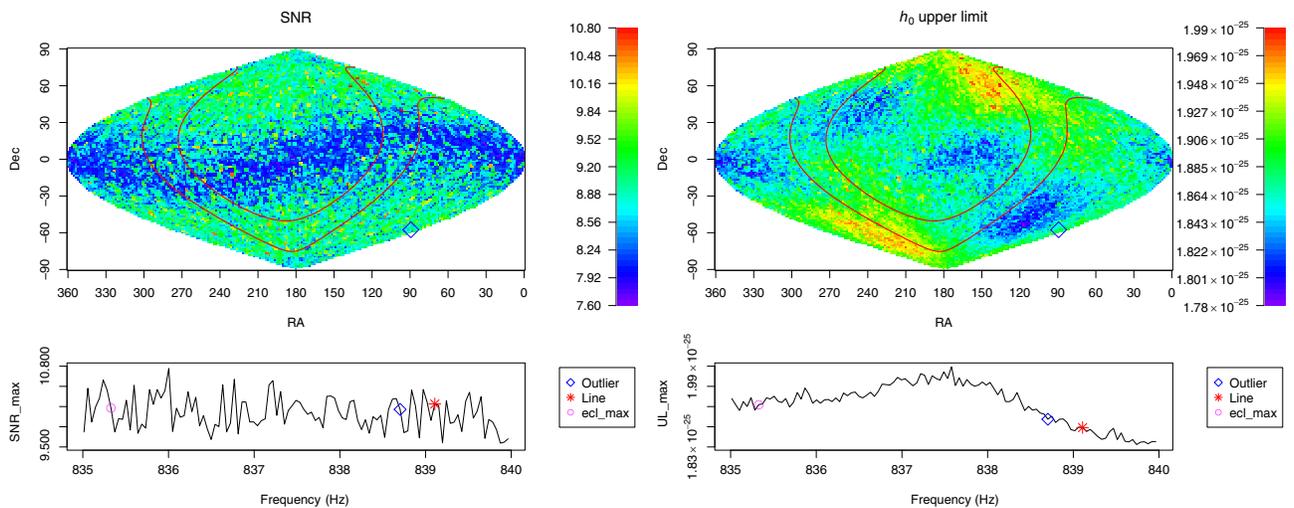


FIG. 4. Summary of atlas data from the bins between 835–840 Hz. The top panels show the highest SNR (left) and upper limit values (right) across the frequency band, for each pixel of the sky map, using equatorial coordinates. The red lines denote the Galactic plane. The blue diamond shows the location of the outlier that is discarded based on the analysis of O3a + b data. The blue band of smaller SNRs near the ecliptic equator is due to large correlations between waveforms of sources in that region. The blue regions in the upper limit plot are due to the lower SNR values in the ecliptic plane, and also occur near the ecliptic poles that are favored by the antenna pattern of the detectors. The bottom panels show the same data as a function of frequency and with the maximum taken over the sky. We mark the frequency of the band where the outlier mentioned above is found, the location of the only known line from the O3 line list in that band, and the band where the maximum SNR is achieved in the ecliptic pole region—a region strongly affected by instrumental lines. The data and code used to produce this plot are available [44]. SNR is achieved in the ecliptic pole region (ecl \times max).

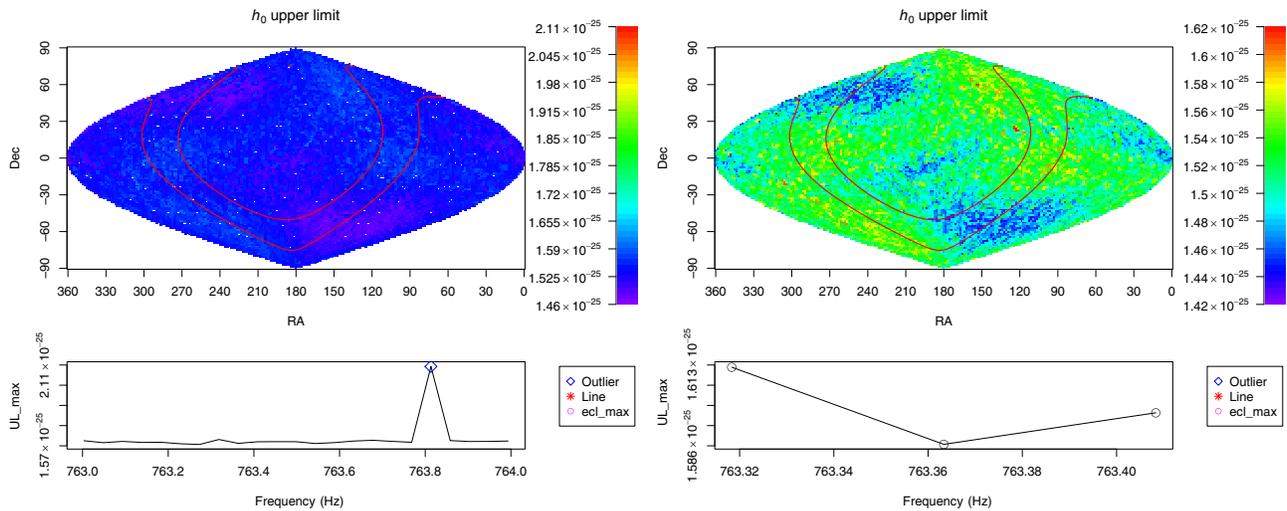


FIG. 5. Both panels here show upper limit results, like the right-hand side plots of Fig. 4. The frequency band is different: 763–764 Hz for the left-hand panels and an enlargement between 763.3 and 763.42 for the right-hand side panels. We choose the 763–764 band to give an idea of what a signal would look like, since the hardware-injected fake signal ip9 falls in this band; see Table III. The signal amplitude is 1.3×10^{-25} and it is lower than the upper limit value at the sky position, as it should be. The enlarged band does not contain the signal parameters and hence it is an example of a “quiet band.”

TABLE II. Outliers surviving the detection pipeline. Outliers marked “ipX” are due to simulated signals “hardware injected” during the science run for validation purposes. Their parameters are listed in Table III. Signal frequencies refer to GPS epoch 1 246 070 000 (2019 Jul 2 02:33:02 UTC).

SNR	f (Hz)	\dot{f} (pHz/s)	RA _{J2000} (deg)	Dec _{J2000} (deg)	Comment
110.4	763.847 32	−0.4	198.886	75.689	Simulated signal ip9
98.9	575.163 51	−0.4	215.255	3.441	Simulated signal ip2
16.7	993.381 80	7.0	282.198	−57.533	Strong narrow band disturbance in H1
16.3	848.919 31	−31.6	39.087	−15.706	Simulated signal ip1
16.1	838.704 99	1.2	10.883	−57.581	Fails O3a + b consistency check

TABLE III. Parameters of the hardware-injected simulated continuous wave signals during the O3 data run—epoch GPS 1 246 070 000 (2019 Jul 2 02:33:02 UTC)—that fall in the parameter range probed by our search.

Label	h_0 10^{-25}	f (Hz)	\dot{f} (pHz/s)	RA _{J2000} (deg)	Dec _{J2000} (deg)
ip1	5.5	848.935	−300	37.39	−29.45
ip2	0.76	575.164	−0.137	215.26	3.44
ip9	1.3	763.847	-1.45×10^{-5}	198.89	75.69

using H1 data alone yields a SNR of 16.3, while more sensitive L1 data give a SNR of only 9.2.

One outlier with SNR just above threshold does not correspond to any identifiable detector artifact; however, a reanalysis of this outlier using public LIGO data from the O3a and O3b [39] runs combined yields a decrease in the SNR—from 16 on O3a data to 13 using O3a + b—with a 6 day coherence length. This is not consistent with what is expected from the IT2 continuous signal model [Eq. (1)] we assume in this search. Thus the outlier is either an astrophysical signal that deviates significantly from the IT2 model or the outlier is due to terrestrial contamination.

ACKNOWLEDGMENTS

The authors thank the scientists, engineers, and technicians of LIGO, whose hard work and dedication produced the data that made this search possible. The search was performed on the ATLAS cluster at AEI Hannover. We thank Bruce Allen, Carsten Aulbert, and Henning Fehrmann for their support. We are grateful to Heinz-Bernd Eggenstein for helpful feedback on the atlas. This research has made use of data or software obtained from the Gravitational Wave Open Science Center, a service of LIGO Laboratory, the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA.

APPENDIX: USING THE FALCON ATLAS TO SEARCH SUPERNOVA REMNANTS

It is easy to use our atlas to examine arbitrary locations on the sky. An example R script (`spatial_index_example2.R`) included with the atlas data shows how to produce upper limits and SNR values for any given sky location. By simply setting right ascension and declination to the coordinates of the Vela Jr or G189.1 + 3.0 supernova remnants, we obtain the data shown in Figs. 6 and 7, respectively.

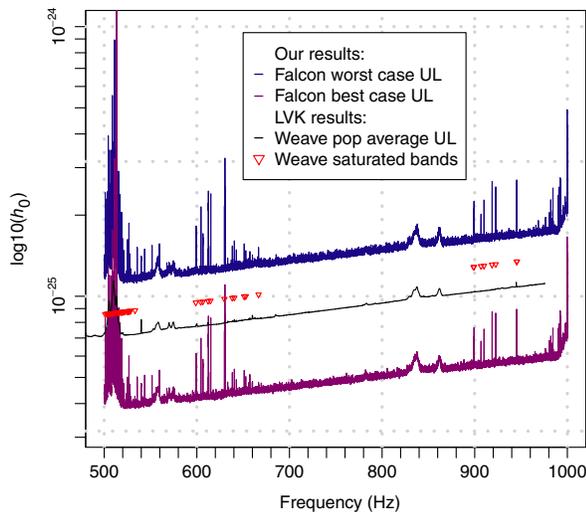


FIG. 6. This plot shows upper limits similar to those shown in Fig. 2, but for the location of Vela Jr. Latest LIGO-Virgo-KAGRA results [50] are shown for comparison. The triangles mark saturated bands for which the Weave results are invalid.

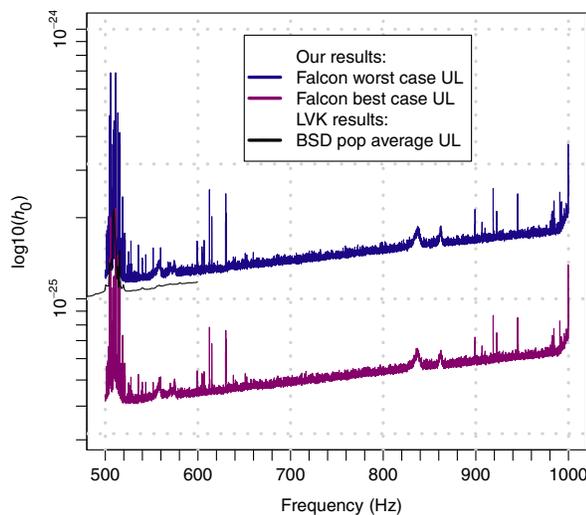


FIG. 7. This plot shows upper limits similar to those shown in Fig. 2, but for the location of G189.1 + 3.0. Latest LIGO-Virgo-KAGRA results [51] are shown for comparison. The LVK upper limit curve was computed as minimum of Hanford, Livingston, and Virgo data.

We also plot the latest LIGO-Virgo-KAGRA results [50,51]. Our spin-down range is smaller than that searched by Refs. [50,51], but we have no excluded bands. Furthermore, our results are rigorous 95% confidence level upper limits, rather than sensitivity estimates.

These plots illustrate that an atlas user can establish rigorous upper limits that improve on the best dedicated directed searches.

These data can be used as is to place limits on the gravitational radiation coming from the direction of these supernova remnants as we have done, or it can be the starting point for a deeper search.

-
- [1] R. Abbott *et al.* (LIGO Scientific Collaboration, VIRGO Collaboration, and KAGRA Collaboration), *GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run*, arXiv:2111.03606.
 - [2] J. Aasi *et al.* (LIGO Scientific Collaboration), *Advanced LIGO*, *Classical Quantum Gravity* **32**, 074001 (2015).
 - [3] F. Acernese *et al.*, *Advanced Virgo: A Second-Generation Interferometric Gravitational Wave Detector*, *Classical Quantum Gravity* **32**, 024001 (2015).
 - [4] H. Abe *et al.* (KAGRA Collaboration), *Performance of the KAGRA Detector During the First Joint Observation with GEO 600 (O3GK)*, *Prog. Theor. Exp. Phys.* (2022) **10.1093/ptep/ptac093**.
 - [5] B. Haskell, M. Priymak, A. Patruno, M. Oppenorth, A. Melatos, and P.D. Lasky, *Detecting Gravitational Waves from Mountains on Neutron Stars in the Advanced Detector Era*, *Mon. Not. R. Astron. Soc.* **450**, 2393 (2015).
 - [6] L. Lindblom, B. J. Owen, and S. M. Morsink, *Gravitational Radiation Instability in Hot Young Neutron Stars*, *Phys. Rev. Lett.* **80**, 4843 (1998).
 - [7] B. J. Owen, *How to Adapt Broad-Band Gravitational-Wave Searches for r -Modes*, *Phys. Rev. D*, **82**, 104002 (2010).
 - [8] A. Arvanitaki and S. Dubovsky, *Exploring the String Axiverse with Precision Black Hole Physics*, *Phys. Rev. D* **83**, 044026 (2011).
 - [9] R. Abbott *et al.* (LIGO Scientific Collaboration, VIRGO Collaboration, and KAGRA Collaboration), *All-Sky Search for Gravitational Wave Emission from Scalar Boson Clouds around Spinning Black Holes in LIGO O3 Data*, *Phys. Rev. D* **105**, 102001 (2022).
 - [10] S. J. Zhu, M. Baryakhtar, M. A. Papa, D. Tsuna, N. Kawanaka, and H. B. Eggenstein, *Characterizing the Continuous Gravitational-Wave Signal from Boson Clouds around Galactic Isolated Black Holes*, *Phys. Rev. D* **102**, 063020 (2020).
 - [11] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, *String Axiverse*, *Phys. Rev. D* **81**, 123530 (2010).
 - [12] K. Riles, *Searches for Continuous-Wave Gravitational Radiation*, arXiv:2206.06447.
 - [13] V. Dergachev, *On Blind Searches for Noise Dominated Signals: A Loosely Coherent Approach*, *Classical Quantum Gravity* **27**, 205017 (2010).

- [14] V. Dergachev, *Loosely Coherent Searches for Sets of Well-Modeled Signals*, *Phys. Rev. D* **85**, 062003 (2012).
- [15] V. Dergachev, *Loosely Coherent Searches for Medium Scale Coherence Lengths*, [arXiv:1807.02351](https://arxiv.org/abs/1807.02351).
- [16] V. Dergachev and M. A. Papa, *Results from the First All-Sky Search for Continuous Gravitational Waves from Small-Ellipticity Sources*, *Phys. Rev. Lett.* **125**, 171101 (2020).
- [17] V. Dergachev and M. A. Papa, *Results from High-Frequency All-Sky Search for Continuous Gravitational Waves from Small-Ellipticity Sources*, *Phys. Rev. D*, **103**, 063019 (2021).
- [18] V. Dergachev and M. A. Papa, *The Search for Continuous Gravitational Waves from Small-Ellipticity Sources at Low Frequencies*, *Phys. Rev. D*, **104**, 043003 (2021).
- [19] 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*, [arXiv:2303.04109](https://arxiv.org/abs/2303.04109).
- [20] Jonathan P. Gardner, John C. Mather, Mark Clampin, Rene Doyon, Matthew A. Greenhouse, Heidi B. Hammel, John B. Hutchings, Peter Jakobsen, Simon J. Lilly, Knox S. Long *et al.*, *The James Webb Space Telescope*, *Space Sci. Rev.* **123**, 485 (2006).
- [21] B. S. Acharya *et al.* (The CTA Consortium), *Science with the Cherenkov Telescope Array* (World Scientific, Singapore, 2019).
- [22] P. E. Dewdney, P. J. Hall, R. T. Schilizzi, and T. J. L. W. Lazio, *The Square Kilometre Array*, *Proc. IEEE* **97**, 1482 (2009).
- [23] M. Aguilar *et al.*, *The Alpha Magnetic Spectrometer (AMS) on the International Space Station: Part II—Results from the First Seven Years*, *Phys. Rep.* **894**, 1 (2021).
- [24] N. M. Law *et al.*, *The Palomar Transient Factory: System Overview, Performance, and First Results*, *Publ. Astron. Soc. Pac.* **121**, 1395 (2009).
- [25] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs, *The Australia Telescope National Facility Pulsar Catalogue*, *Astron. J.* **129**, 1993 (2005).
- [26] M. Wenger, F. Ochsenbein, D. Egret, P. Dubois, F. Bonnarel, S. Borde, F. Genova, G. Jasniewicz, S. Laloč, S. Lesteven, and R. Monier, *The SIMBAD Astronomical Database—The CDS Reference Database for Astronomical Objects*, *Astron. Astrophys. Suppl. Ser.* **143**, 9 (2000).
- [27] Mappable Vector Library, <https://github.com/volodya31415/libMVL>.
- [28] R Package for Mappable Vector Library, <https://cran.r-project.org/package=RMVL>.
- [29] N. K. Johnson-McDaniel and B. J. Owen, *Maximum Elastic Deformations of Relativistic Stars*, *Phys. Rev. D* **88**, 044004 (2013).
- [30] F. Gittins, N. Andersson, and D. I. Jones, *Modelling Neutron Star Mountains*, *Mon. Not. R. Astron. Soc.* **500**, 5570 (2020).
- [31] F. Gittins and N. Andersson, *Modelling Neutron Star Mountains in Relativity*, *Mon. Not. R. Astron. Soc.* **507**, 116 (2021).
- [32] G. Pagliaro, M. A. Papa, J. Ming, J. Lian, D. Tsuna, C. Maraston, and D. Thomas, *Continuous Gravitational Waves from Galactic Neutron Stars: Demography, Detectability and Prospects*, [arXiv:2303.04714](https://arxiv.org/abs/2303.04714).
- [33] G. Woan, M. D. Pitkin, B. Haskell, D. I. Jones, and P. D. Lasky, *Evidence for a Minimum Ellipticity in Millisecond Pulsars*, *Astrophys. J. Lett.* **863**, L40 (2018).
- [34] P. Jaranowski, A. Krolak, and B. F. Schutz, *Data Analysis of Gravitational-Wave Signals from Spinning Neutron Stars: The Signal and Its Detection*, *Phys. Rev. D* **58**, 063001 (1998).
- [35] S. R. Valluri, V. Dergachev, X. Zhang, and F. A. Chishtie, *Fourier Transform of the Continuous Gravitational Wave Signal*, *Phys. Rev. D* **104**, 024065 (2021).
- [36] V. Dergachev, *A Novel Universal Statistic for Computing Upper Limits in Ill-Behaved Background*, *Phys. Rev. D* **87**, 062001 (2013).
- [37] ATLAS Computing Cluster, <https://www.aei.mpg.de/43564/atlas-computing-cluster>.
- [38] The O3a Data Release, [10.7935/nfnt-hm34](https://arxiv.org/abs/10.7935/nfnt-hm34).
- [39] The O3b Data Release, [10.7935/pr1e-j706](https://arxiv.org/abs/10.7935/pr1e-j706).
- [40] 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).
- [41] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), *All-Sky Search for Continuous Gravitational Waves from Isolated Neutron Stars in the Early O3 LIGO Data*, *Phys. Rev. D*, **104**, 082004 (2021).
- [42] D. Davis *et al.*, *LIGO Detector Characterization in the Second and Third Observing Runs*, *Classical Quantum Gravity* **38**, 135014 (2021).
- [43] B. Steltner, M. A. Papa, and H. B. Eggenstein, *Identification and Removal of Non-Gaussian Noise Transients for Gravitational-Wave Searches*, *Phys. Rev. D* **105**, 022005 (2022).
- [44] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevX.13.021020> for numerical values of upper limits.
- [45] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), *All-Sky Search for Continuous Gravitational Waves from Isolated Neutron Stars Using Advanced LIGO and Advanced Virgo O3 Data*, *Phys. Rev. D* **106**, 102008 (2022).
- [46] P. B. Covas, M. A. Papa, R. Prix, and B. J. Owen, *Constraints on r -Modes and Mountains on Millisecond Neutron Stars in Binary Systems*, *Astrophys. J. Lett.* **929**, L19 (2022).
- [47] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), *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).
- [48] A. L. Miller, S. Clesse, F. De Lillo, G. Bruno, A. Depasse, and A. Tanasijczuk, *Probing Planetary-Mass Primordial Black Holes with Continuous Gravitational Waves*, *Phys. Dark Universe* **32**, 100836 (2021).
- [49] C. Biwer, D. Barker, J. C. Batch, J. Betzwieser, R. P. Fisher, E. Goetz, S. Kandhasamy, S. Karki, J. S. Kissel, A. P. Lundgren *et al.*, *Validating Gravitational-Wave Detections: The Advanced LIGO Hardware Injection System*, *Phys. Rev. D* **95**, 062002 (2017).

- [50] R. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Search of the Early O3 LIGO Data for Continuous Gravitational Waves from the Cassiopeia A and Vela Jr. Supernova Remnants*, *Phys. Rev. D* **105**, 082005 (2022).
- [51] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), *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).