



Deep Einstein@Home All-sky Search for Continuous Gravitational Waves in LIGO O3 Public Data

B. Steltner^{1,2} , M. A. Papa^{1,2} , H.-B. Eggenstein^{1,2} , R. Prix^{1,2}, M. Bensch^{1,2}, B. Allen^{1,2} , and B. Machenschalk^{1,2}

¹ Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Callinstrasse 38, D-30167 Hannover, Germany; benjamin.steltner@aei.mpg.de, maria.alessandra.papa@aei.mpg.de

² Leibniz Universität Hannover, D-30167 Hannover, Germany

Received 2023 March 10; revised 2023 May 31; accepted 2023 May 31; published 2023 July 18

Abstract

We present the results of an all-sky search for continuous gravitational waves in the public LIGO O3 data. The search covers signal frequencies $20.0 \text{ Hz} \leq f \leq 800.0 \text{ Hz}$ and a spin-down range down to $-2.6 \times 10^{-9} \text{ Hz s}^{-1}$, motivated by detectability studies on synthetic populations of Galactic neutron stars. This search is the most sensitive all-sky search to date in this frequency/spin-down region. The initial search was performed using the first half of the public LIGO O3 data (O3a), utilizing graphical processing units provided in equal parts by the volunteers of the Einstein@Home computing project and by the ATLAS cluster. After a hierarchical follow-up in seven stages, 12 candidates remain. Six are discarded at the eighth stage, by using the remaining O3 LIGO data (O3b). The surviving six can be ascribed to continuous-wave fake signals present in the LIGO data for validation purposes. We recover these fake signals with very high accuracy with our last stage search, which coherently combines all O3 data. Based on our results, we set upper limits on the gravitational-wave amplitude h_0 and translate these into upper limits on the neutron star ellipticity and on the r -mode amplitude. The most stringent upper limits are at 203 Hz, with $h_0 = 8.1 \times 10^{-26}$ at the 90% confidence level. Our results exclude isolated neutron stars rotating faster than 5 ms with ellipticities greater than $5 \times 10^{-8} \left[\frac{d}{100 \text{ pc}} \right]$ within a distance d from Earth and r -mode amplitudes $\alpha \geq 10^{-5} \left[\frac{d}{100 \text{ pc}} \right]$ for neutron stars spinning faster than 150 Hz.

Unified Astronomy Thesaurus concepts: Gravitational waves (678); Gravitational wave astronomy (675); Neutron stars (1108)

Supporting material: data behind figure, tar.gz file

1. Introduction

Continuous gravitational waves are nearly monochromatic, long-lasting signals. They could come from fast-rotating nonaxisymmetric neutron stars, from the excitation of unstable r -modes (Owen et al. 1998; Lasky 2015), the fast inspiral of dark matter objects (Horowitz & Reddy 2019; Horowitz et al. 2020), or superradiant emission of axion-like particles around black holes (Arvanitaki et al. 2015; Zhu et al. 2020).

The detection of a continuous gravitational wave is still elusive. Compared to the already detected gravitational waves of compact binary coalescences (Abbott et al. 2019, 2021a, 2021b; Nitz et al. 2019, 2020, 2021a, 2021b; Venumadhav et al. 2019, 2020; Olsen et al. 2022), the continuous gravitational-wave amplitude at Earth is orders of magnitude smaller. However, since the signal is long-lasting, one can integrate it over many months and increase the signal-to-noise ratio.

When the waveform parameters are not known, broad parameter searches are carried out, and they are expensive because the number of waveforms that can be resolved over many months of observational data is extremely large.

In this paper, we present an all-sky search for unknown, isolated neutron stars with a gravitational-wave frequency $20.0 \text{ Hz} \leq f \leq 800.0 \text{ Hz}$ and spin-down $-2.6 \times 10^{-9} \text{ Hz s}^{-1}$

$\leq \dot{f} \leq 2.6 \times 10^{-10} \text{ Hz s}^{-1}$, carried out on the distributed computing volunteer project Einstein@Home and the ATLAS supercomputer at the Max Planck Institute (MPI) for Gravitational Physics in Hannover. The frequency–spin-down range is based on the predictions of Pagliaro et al. (2023), according to which more than 95% of the potentially detectable sources lie in this range.

We use the public data of the third observing run (O3) of the two Advanced LIGO detectors, near Hanford (LHO) and Livingston (LLO), respectively (Aasi et al. 2015; Abbott et al. 2021c). Since continuous-wave searches are computationally limited, the investment of computing resources in the processing of an additional data stream has to be carefully weighted against the gains in sensitivity from it. Our investigations advise against including O3 data from the Virgo detector in this search, due to its lower sensitivity (Acernese et al. 2015; Abbott et al. 2023).

We use a staged approach: we search half of the data (O3a) and keep the other half (O3b) to verify any candidate that survives the first search. The O3a search is actually a hierarchy of seven stages, beginning with a computationally intensive step, which is also farmed out on the volunteer computing project Einstein@Home. Finally, any candidate surviving the O3b stage (Stage 8) is confirmed with a fully coherent search on the entire data set, O3a+b, based on which the signal parameters are most accurately estimated.

The plan of the paper is as follows. Section 2 describes the signal model and Section 3 the search methodology. The Einstein@Home search is described in Section 4; the



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

hierarchical follow-ups in Section 5. Results are presented in Section 6 and conclusions in Section 7.

2. The Signal

The waveforms $h(t)$ that we target in this search are fairly simple: nearly monochromatic signals with frequency and amplitude modulation due to the Earth’s motion. At the gravitational-wave detector, they take the form (Jaranowski et al. 1998)

$$h(t) = F_+(\alpha, \delta, \psi; t)h_+(t) + F_\times(\alpha, \delta, \psi; t)h_\times(t), \quad (1)$$

where $F_+(\alpha, \delta, \psi; t)$ and $F_\times(\alpha, \delta, \psi; t)$ are the detector beam pattern functions for the “+” and “ $\times$ ” polarizations, (α, δ) are the right ascension and declination of the source, ψ is the polarization angle, and t is the time at the detector. The waveforms $h_+(t)$ and $h_\times(t)$ take the form

$$\begin{aligned} h_+(t) &= A_+ \cos \Phi(t) \\ h_\times(t) &= A_\times \sin \Phi(t), \end{aligned} \quad (2)$$

with the “+” and “ $\times$ ” amplitudes:

$$\begin{aligned} A_+ &= \frac{1}{2}h_0(1 + \cos^2 \iota) \\ A_\times &= h_0 \cos \iota. \end{aligned} \quad (3)$$

$h_0 \geq 0$ is the intrinsic gravitational-wave amplitude, $0 \leq \iota \leq \pi$ is the angle between the total angular momentum of the star and the line of sight, and $\Phi(t)$ is the phase of the gravitational-wave signal at the time t . If τ_{SSB} is the arrival time of the wave with phase $\Phi(t)$ at the solar system barycenter, then $\Phi(t) = \Phi(\tau_{\text{SSB}}(t))$. The gravitational-wave phase as a function of τ_{SSB} is assumed to be

$$\begin{aligned} \Phi(\tau_{\text{SSB}}) &= \Phi_0 + 2\pi[f(\tau_{\text{SSB}} - \tau_{0\text{SSB}}) \\ &\quad + \frac{1}{2}\dot{f}(\tau_{\text{SSB}} - \tau_{0\text{SSB}})^2]. \end{aligned} \quad (4)$$

We take $\tau_{0\text{SSB}} = 1246070525.0$ (Barycentric Dynamical Time in GPS seconds) as a reference time.

We assume that in our target population, the following quantities are uniformly distributed: $20.0 \text{ Hz} \leq f \leq 800.0 \text{ Hz}$, $|\cos \iota| \leq 1$, $|\psi| \leq \pi/4$, source position $0 \leq \alpha < 2\pi$, and $|\sin \delta| \leq 1$, each distributed uniformly. We assume that the spin-down is distributed log-uniformly in our search range, reflecting our ignorance of the actual spin-down distribution.

As we will see, various parameters pertaining to the multistage search presented in this paper are set based on the recovery rate of the search performed over the same reference signal population. Now we describe this reference population. We use ≈ 1600 signals. The frequency, spin-down, position, $\cos \iota$, and ψ parameters are distributed as described above. The amplitudes are such that the sensitivity depth (defined in Equation (7)) is uniformly distributed in $\mathcal{D} \in [50, 65] [1/\sqrt{\text{Hz}}]$, bracketing a competitive but realistic sensitivity depth value of $56 [1/\sqrt{\text{Hz}}]$.

3. Generalities of the Searches

3.1. The Data

We use O3 calibrated data with linear and nonlinear noise subtraction (Davis et al. 2019; Vajente et al. 2020), which




Figure 1. Segmentation for the various searches and input data (SFTs in LHO and LLO). Stages 7, 8, and 9 are fully coherent, using O3a, O3b, and O3a+b, respectively.

removes spurious noise due to laser beam jitter, calibration lines, power mains, and background noise.

As with previous Einstein@Home searches, we remove noise negatively affecting our search, namely *lines* in the frequency domain and *glitches* in the time domain, as described in Steltner et al. (2022b).

The data are split into our usual format of short-time-baseline Fourier transforms (SFTs), with a half-hour time baseline. These are grouped into segments of variable duration for the different coherence times employed in the follow-up, as shown in Figure 1.

3.2. The Search

We utilize two detection statistics, the $\mathcal{F}$ -statistic (Jaranowski et al. 1998; Cutler & Schutz 2005) and the line- and transient-line-robust statistic $\hat{\beta}_{\text{S/GLtL}}$ (Keitel 2016). The data are split into N_{seg} segments of equal span T_{coh} . The data of both detectors in each segment are combined coherently and the detection statistic values for each segment are calculated. The $\mathcal{F}$ -statistic values from the N_{seg} segments are summed to yield the final semicoherent detection statistic:

$$\hat{\mathcal{F}}(x, \lambda_t) = \frac{1}{N_{\text{seg}}} \sum_{i=1}^{N_{\text{seg}}} \mathcal{F}_i(x, \lambda_t), \quad (5)$$

where x indicates the data and λ_t the parameters defining the template waveform. For an isolated rotating neutron star, the template waveform is defined by the gravitational-wave frequency f , its derivative $\dot{f}$ (often spin-down), and the sky position (α, δ) : $\lambda = (f, \dot{f}, \alpha, \delta)$.

The $\hat{\mathcal{F}}$ -statistic is computed from the log-likelihood ratio of the signal hypothesis to the Gaussian noise hypothesis, whereas the $\hat{\beta}_{\text{S/GLtL}}$ -statistic tests the signal hypothesis against an expanded noise hypothesis, i.e., “G” Gaussian noise or “L” lines or “tL” transient lines (Keitel 2016). So, while the $\mathcal{F}$ -statistic is susceptible to disturbances generated by spectral lines, the $\hat{\beta}_{\text{S/GLtL}}$ -statistic greatly reduces the number of

Table 1
Overview of the Full Hierarchy of Searches

Search	T_{coh} (hr)	N_{seg}	δf (μHz)	$\delta \dot{f}$ ($10^{-14} \text{ Hz s}^{-1}$)	m_{sky}	$\langle \mu \rangle$ (10^{-2})	Δf (μHz)	$\Delta \dot{f}$ ($10^{-14} \text{ Hz s}^{-1}$)	$\frac{r_{\text{sky}}}{d(0.002)}$	N_{in}	N_{out}
Stage 0	120	37	2	60	0.002	56	Full range	Full range	All sky	6.7×10^{18}	3513855
Stage 1	120	37	1	15	0.0002	31	1000	11,250	10.0	3,513,855	386,429
Stage 2	120	37	1	2	2×10^{-6}	2.5	50	1200	1.0	386,429	35,635
Stage 3	240	19	0.1	0.2	1.0×10^{-8}	0.06	5	200	0.2	35,635	5116
Stage 4	490	9	0.013	0.064	3.2×10^{-10}	<0.01	0.5	35	0.04	5116	1387
Stage 5	1100	4	0.001	0.032	1.4×10^{-10}	<0.01	0.2	20	0.017	1387	310
Stage 6	2200	2	0.001	0.009	2.8×10^{-11}	<0.01	0.1	8.5	0.008	310	54
Stage 7	Coherent	1	0.001	0.0063	1×10^{-11}	<0.01	0.07	6	0.0057	54	12
Stage 8	O3b coh.	1	0.001	0.0063	1×10^{-11}	<0.01	$\geq 0.07^a$	6	0.0057	12	6
Stage 9 ^b	O3a+b coh.	1	0.001	0.0063	1×10^{-11}	<0.01	$\geq 0.07^a$	6	0.0057	6	6

Notes. The columns show the values of the following parameters. Column (1): the stage number. Column (2): the coherent time baseline T_{coh} of each segment. Column (3): the number of segments N_{seg} . Columns (4)–(6): the grid spacings δf , $\delta \dot{f}$, and m_{sky} . Column (7): the average mismatch $\langle \mu \rangle$. Columns (8)–(10): the parameter space volume searched around each candidate, Δf , $\Delta \dot{f}$, and r_{sky} . The radius r_{sky} is expressed in units of the side of the hexagonal sky-grid tile of the Stage 0 search (Equation (6)). Columns (11)–(12): the number of candidates searched (N_{in}) and how many of those survive and make it to the next stage (N_{out}).

^a Since the reference time of the O3b search is different than the reference time of the O3a stages, an uncertainty in spin-down value produces an uncertainty in frequency. The nominal 0.07 μHz value would only hold if the signal spin-down were precisely known.

^b Since even at the previous stage the only surviving candidates are the hardware injections, we carry out this stage to demonstrate the accuracy in signal recovery.

candidates from these disturbances. Therefore, we rank the results using the latter.

For efficiency reasons, the detection statistic value is first computed on a coarse template grid, and then approximated on a finer grid (Pletsch 2010). At the end, the detection statistic of the highest-ranking results is recomputed exactly at the fine-grid template point. The recomputed quantities are indicated by a subscript r . These are the results returned to the Einstein@Home central server. We refer to the waveform templates and the associated detection statistic values of the returned results as *candidates*.

The grid spacings are chosen to minimize the computational cost for a given average loss in the detection statistic due to signal/template mismatch—this quantity is known as the “average mismatch” and is indicated with $\langle \mu \rangle$. The average mismatch value is chosen based on computational feasibility. Table 1 shows the spacings and the average mismatch for all stages. The first search, Stage 0, is the most challenging, because the computational cost of surveying the entire parameter space is very high, and this results in a grid with a high mismatch—about 56%. Figure 2 shows the number of templates in this grid, as a function of frequency. The total number of templates searched is 6.7×10^{18} . The total number of coarse-grid templates is $\approx 2.7 \times 10^{16}$.

The grids in frequency and spin-down are defined by δf and $\delta \dot{f}$, respectively, and these do not change across the search range. Conversely, the sky grid varies with frequency, becoming finer at higher frequencies. Our sky grids are approximately uniform on the celestial sphere orthogonally projected on the ecliptic plane and are defined by the parameter m_{sky} . For the equations defining the projected coordinates, see Equations (14) and (15) in Singh et al. (2017). The tiling is a hexagonal covering of the unit circle with each hexagon having the edge length d :

$$d(m_{\text{sky}}) = 0.15 \sqrt{m_{\text{sky}}} \left[\frac{100 \text{ Hz}}{f} \right]. \quad (6)$$




Figure 2. Number of templates searched, per 50 mHz band and cumulative, by this search and by the Einstein@Home O2 all-sky search (Steltner et al. 2021) ending at 585.15 Hz. The sky resolution increases with frequency (see Equation (6)), and so does the number of templates. This search uses a finer resolution than Steltner et al. (2021) for the same range in spin-down and sky, thus requiring more templates overall.

4. Stage 0: The First Search

4.1. The Distribution of the Computational Load

The first eight searches use LIGO data from the first half of the third observing run (O3a), i.e., between GPS time 1238166018 (2019 April 1 15:00:00 GMT) and 1254150018 (2019 October 3 15:00:00 GMT).

All stages employ the ATLAS cluster. Stage 0 additionally leverages the computing power of the Einstein@Home project. Einstein@Home is a distributed volunteer computing project, built upon the BOINC infrastructure³ (Anderson 2004; Anderson et al. 2006), where volunteers can spend their idle computational resources to solve scientific problems that

³ <https://boinc.berkeley.edu>

require large amounts of computing power. ATLAS is the supercomputer cluster at the MPI for Gravitational Physics in Hannover.⁴

This is the first Einstein@Home continuous gravitational-wave search primarily run on graphical processing units (GPUs). The advantage of using GPUs is that certain instructions can be efficiently parallelized, improving the performance by more than an order of magnitude, compared to CPUs.

The search is split into work units (WUs), which run on an average GPU for 10–30 minutes. A total of 55.7 million WUs were computed, totaling over 2000 yr of computing on a single GPU. Each WU searches 1.2×10^{11} template waveforms, corresponding to half a hertz in frequency, the full spin-down range, and a portion of the sky, and returns a so-called “top-list” containing the top-ranking 30,000 results.

A fraction of the top-list results returned to the server are considered for further processing. In general, the more results that are considered, the lower is the smallest detectable signal. For every search, we consider as many results as we possibly can, given computational constraints.

The overall number of top-list results is 1.7×10^{12} , which effectively is about a factor of ≈ 10 higher compared to the previous all-sky Einstein@Home search (Steltner et al. 2021). This is consistent with the fact that we search more templates compared to Steltner et al. (2021), as Figure 2 shows.

The fraction of the top-list results that we consider for further processing is comparable to our previous search. But since the number of top-list results is ≈ 10 times larger than our previous search, we have to process ≈ 10 times more results. The first step in processing so many more results is enabled by a new and enormously more efficient clustering method (Steltner et al. 2022a).

Despite the significant achievements of the cleaning efforts, there are still disturbances in the data, which, if loud enough, can saturate the entire half-Hz top-list and render it useless. Such disturbances are however typically concentrated in a frequency band much smaller than half a hertz, so to avoid them saturating the entire half-Hz top-list, the search code on the volunteer host maintains 10 *independent* top-lists, one for each of the ten 50 mHz subbands in the half-Hz band. The final half-Hz top-list is the union of all 10 of the 50 mHz top-lists. Each of 50 mHz top-list comprises 3000 results.

4.2. Post-processing

The following post-processing steps are performed on the results returned from the first search:

Banding: all results for each 50 mHz band—from the full spin-down range and all-sky points—are gathered. Based on them, a series of diagnostics are produced, which help identify bands affected by disturbances (see Section IIIB in Abbott et al. 2017). We find that $\approx 1.3\%$ of the 50 mHz bands are disturbed, but they contribute 20% of the candidates. As explained in Section 6.2, these bands will be excluded from the upper limit statements, but candidates from these bands will in general be followed up.

Clustering: since our search grids are somewhat oversampled to reduce the signal-to-template waveform mismatch,

nearby templates are not independent. Hence, a disturbance or a signal produces some to many nearby results, while statistical fluctuations rarely “clump.” Our clustering method identifies results that are due to the same root cause, so that they can be considered as a single candidate.

Every cluster is identified by “a seed,” i.e., a set of signal parameters λ_{seed} , and by an uncertainty range $\Delta\lambda$. The signal parameters are different for every cluster, whereas the uncertainty range is exactly the same for all clusters. The meaning of the uncertainty range $\Delta\lambda$ is the following: $>99.9\%$ of signals of the reference population described at the end of Section 2 give rise, after Stage 0, to at least a cluster whose seed parameters are within a distance $\Delta\lambda$ of the signal parameters.

We indicate the uncertainty intervals with Δf , $\Delta \dot{f}$, and r_{sky} . The uncertainty region in the sky is a circle in the orthogonally projected ecliptic plane centered at the candidates’ sky position, with radius r_{sky} .

The clustering parameters are determined based on search-and-recoveries on the reference signal population after Stage 0. A clustering setup is chosen that minimizes the amplitude of the weakest signals recovered, while the number of false alarms remains below a given threshold, determined by the total amount of time we want to devote to the Stage 1 follow-up. With a Stage 1 follow-up of a few weeks, the 90% recovery rate of the chosen clustering setup corresponds to a population of signals with amplitude h_0 , such that $\mathcal{D} = 56 [1/\sqrt{\text{Hz}}]$. For a fake signal to be counted as recovered, there needs to exist a seed that can be associated with the injection, which was not there in the data without the injection. We remind the reader that the sensitivity depth $\mathcal{D}$, first introduced by Behnke et al. (2015), is defined as

$$\mathcal{D} = \sqrt{S_n(f)} / h_0(f) [1/\sqrt{\text{Hz}}], \quad (7)$$

with $h_0(f)$ being the continuous gravitational-wave intrinsic strain amplitude. If h_0 is the upper limit from a search at frequency f , then $\mathcal{D}$ describes the sensitivity of that search in terms of “how deep” below the noise level S_n the search can detect signals. But Equation (7) can also be seen as defining the amplitudes of the population of signals at different frequencies, which would be $\approx$ equally well detected by a given search pipeline: fix the detection pipeline, this determines the value of $\mathcal{D}(f)$, and Equation (7) gives the amplitude $h_0(f)$ of the smallest detectable signal. This is used in search-and-recovery simulations that aim at characterizing the detection efficiency of a pipeline, or a piece thereof, as done above for the clustering.

The clustering reduces the 1.7×10^{12} results to a more manageable set of ≈ 3.5 million seeds. For the remainder of the paper, we may also refer to the cluster seeds as the “Stage 0 candidates” or simply as “candidates.”

Follow-up: we follow up the 3.5 million candidates as detailed in the next section. The average number of candidates per 50 mHz band ranges between ≈ 50 and 350 candidates, as shown in Figure 3. We do not follow up any candidate from 50 mHz bands with more than 10,000 candidates, as this is a clear indication that the band is affected by disturbances and the candidates in it are due to the disturbances. 57 bands are hence excluded from the follow-up and they are provided in the supplemental .tar.gz

⁴ <https://www.atlas.aei.uni-hannover.de/>




Figure 3. Stage 0 candidates from every 50 MHz band. Candidates from bands with more than 10,000 candidates per band are not followed up; the rest are.

package. These bands are also excluded from the upper limit statements.

From Figure 3, we see that four of the seven fake signals added to the data for validation—the so-called hardware injection signals—are in relatively “quiet” bands (the ones at higher frequencies), and three (at ~ 26.3 , 31.4 , and 52.8 Hz) are instead in bands that are clearly affected by some excess. The 52.8 Hz injection is very loud and it is solely responsible for the excess. The 31.4 Hz injection produces a very weak signal in the search results—in fact, this is the injection that we are not able to detect—and the excess is due to a disturbance. The 26.3 Hz injection is detected, but the excess comes from a disturbance also present in the band. Figure 4 shows the search results in these three bands and illustrates these three different situations. More information on the hardware injections is given in Section 6.1.

An overview of all Stage 0 search results is given in Figure 5.

5. The Follow-up Searches

A hierarchical follow-up of the clustered candidates from Stage 0 is performed.

An uncertainty region can be defined for each stage, in the same way as done for clustering after Stage 0 (see the bullet point “Clustering” in Section 4.2).

At each stage n , the uncertainty region from stage $n - 1$ around each surviving candidate is searched.

From one stage to the next, the grid resolution is increased. For Stages 3–7, the coherence time T_{coh} is $\approx$ doubled, whereas for Stages 1–2 it stays the same as for Stage 0. The reason is that we would not have enough computing power to usefully follow up all Stage 0–1 candidates with a search having a $T_{\text{coh}} > 120$ hr: the parameter space regions searched in the early stages are much larger than those searched in the later stages and the computational cost per unit volume of parameter space steeply increases with the coherence time T_{coh} (Brady & Creighton 2000). So in the first stages, we keep the computational cost at bay by not increasing T_{coh} . We are however still able to increase the sensitivity of the search by decreasing the grid mismatch, i.e., by using a finer grid. This in turn decreases the uncertainty on the candidate parameters, shrinking the search volume of the next stage. Starting with Stage 3, the search volume is small enough that it becomes computationally feasible to $\approx$ double the coherence time T_{coh} in each stage.

Stages 7, 8, and 9 are fully coherent on O3a, O3b, and O3a + b data, respectively. The search setups and covered regions are the same, as explained in Section 5.1.

The parameters defining the search setups and searched/surviving candidates are all given in Table 1. The mismatch μ distributions are shown in Figure 6.

Only the highest detection statistic result from each follow-up is considered, and that becomes the new representative candidate for that stage. For stages $a = 1, 2, 3$, a candidate is vetoed unless both its $2\mathcal{F}_r$ and $\hat{\beta}_{\text{S/GLIL},r}$ values exceed the thresholds given in Table 2. For stages $a = 4, \dots, 8$, candidates are vetoed unless their

$$R^a = \frac{2\mathcal{F}_r^{\text{Stage } a} - 2\mathcal{F}_r^{\text{Stage } 0}}{2\mathcal{F}_r^{\text{Stage } 0} - 4} \quad (8)$$

exceeds the $R^{a-\text{thr}}$ threshold value given in Table 2.

The $2\mathcal{F}_r^{\text{thr}}$, $\hat{\beta}_{\text{S/GLIL},r}^{\text{thr}}$, and $R^{a-\text{thr}}$ are determined by adding fake signals from our target population (see Section 2) and searching for them, exactly as done in Stage 0 and in the following stages. The total number of fake signals is ≈ 1600 and the thresholds are set so that none of the signals are discarded by the vetoes, corresponding to a false dismissal of $< 99.9\%$. The results are shown in the top right-hand-side plots of Figure 7.

Of the 3.5 million candidates that are fed to the hierarchical follow-up, only 12 remain after Stage 7.

5.1. Follow-up of the 12 Remaining Candidates

We follow up the 12 candidates that survive the seven follow-up stages utilizing a different data set, namely the second half of O3 (O3b). Our data span 147 days, starting at GPS time 1256655667 (2019 November 1 15:00:49 GMT) and ending at GPS time 1269363493 (2020 March 27 16:57:55 GMT). We perform a fully coherent search with the same grid spacings as the fully coherent search on O3a data. Since O3b spans a period of time about ≈ 36 days shorter than O3a, this results in a smaller average mismatch. However, due to the smaller amount of data and due to a slightly worse detector sensitivity, all in all, the search sensitivity using O3b data is reduced with respect to O3a by about 15%. This can be seen in the two bottom panels of Figure 7, comparing the respective signal results, and it yields a lower value of $R^{8-\text{thr}}$.

The uncertainty ranges around the candidates surviving Stages 7 and 8, are larger (by 10%–30%) than the uncertainty regions at Stage 6. We hence center the Stage 8 and 9 follow-ups around the corresponding Stage 6 candidates and use the Stage 6 uncertainty region. The Stage 8 and 9 Monte Carlo simulations are of course performed consistently, i.e., based on the Stage 6 candidates surviving Stage 7. The increase in uncertainty for Stages 7 and 8 is somewhat unexpected, and even though we were able to verify it in a number of ways, we do not fully understand its origin.

After the follow-up on O3b data (Stage 8), only six candidates survive, and they are all due to hardware injections.

6. Results

6.1. Recovery of the Hardware Injections

The hardware injections are signals added to the data by directly moving the detector mirrors in order to provide a check of the entire detection chain. These are de facto reference




Figure 4. Frequency–spin-down plots for the three low-frequency hardware injections with IDs 10 (left), 11 (middle), and 5 (right; LIGO & Virgo 2022). Color-coded are the detection statistic of search results within 50 mHz of the injection parameters, the full spin-down range, and the whole sky. The black crosses mark candidates that are followed up. The red circle indicates the hardware injection parameters. The plots showcase three interesting cases. Left: a relatively loud hardware injection is recovered in the vicinity of a disturbance. Middle: this shows the only not-recovered injection (ID 11). The larger detection statistic values and candidates in this band have nothing to do with the injection signal, but come from a disturbance. Right: the very loud hardware injection (ID 5) also leads to enhanced detection statistic values and candidates at nearby parameter points.




Figure 5. The detection statistic values in $2\overline{\mathcal{F}}_r$ and $\hat{\beta}_{S/GLTLr}$ of the loudest candidates in each 50 mHz band (gray triangles, which appear black when many of them are superimposed upon one another) and of each candidate selected for follow-up (circles). This shows that the candidates that we are following up have $2\overline{\mathcal{F}}_r$ values up to 40% smaller than if only the loudest candidates per 50 mHz band had been selected. The hardware injections are denoted in green. The 3σ interval around the expected highest $2\overline{\mathcal{F}}_r$, determined by the number of searched templates is given in orange and is in good agreement with our data. We also mark (red circles) follow-up candidates from disturbed bands (see the bullet point “Banding” in Section 4.2).

signals that serve as standard detection benchmarks for any continuous-wave search pipeline.

Seven hardware injections fall in our search range, specifically those with IDs 0, 2, 3, 5, 9, 10, and 11 (LIGO & Virgo 2022). We recover all but one.

The missed hardware injection has ID = 11, and it is at ~ 31.4 Hz. Its strain amplitude lies just below our upper limit, but its inclination is not particularly unfavorable. The reason why it is not detected is that its parameter values lie at a high mismatch point within the grid, and the resulting detection statistic value is low enough that in the original Stage 0 results there is no candidate associated with this hardware injection.

The O3a+b candidates associated with the hardware injections are within $\pm 10^{-9}$ Hz, $\pm 10^{-15}$ Hz s $^{-1}$, and $0.3''$ in sky (see Table 3) of the correct parameter values. This remarkable accuracy is expected from long-baseline

observations, and it is one of the promises of the science of continuous gravitational waves.

The O3a+b search setup is not optimal; rather, it is a practical search with improved sensitivity with respect to all the previous stages. This means that probably even higher parameter estimation accuracy could be obtained.

Three of the candidates surviving Stage 7 are “secondaries” associated to hardware injections 2, 9, and 10. Compared to the primaries, they are not as significant and lie at a much greater distance from the true signal values: $> 2 \times 10^{-4}$ Hz in frequency and $> 6 \times 10^{-11}$ Hz s $^{-1}$ in spin-down. None of the secondaries survive Stage 8.

6.2. Upper Limits

Based on our null results, we set 90% confidence frequentist upper limits on the detectable intrinsic gravitational-wave




Figure 6. Mismatch distributions of all stages based on 1000 injection-and-recovery Monte Carlo simulations. Due to the high cost of Stage 0, it features a large mismatch of $\sim 56\%$. The setups of later stages are chosen so that the detection statistics increase from one stage to the next, and this is achieved by increasing T_{coh} and/or decreasing the mismatch. The finer setups become possible by the shrinking of the uncertainty region and by the decreasing number of candidates per stage.

Table 2

The Threshold Value R^a Used to Veto Candidates of Stage $a > 3$ (Equation (8)); the Absolute Thresholds Used in Stages 1–3: $2\bar{\mathcal{F}}_r^{\text{thr}}$, $\hat{\beta}_{\text{S/GLILr}}^{\text{thr}}$

Search	$R^{a-\text{thr}}$	$2\bar{\mathcal{F}}_r^{\text{thr}}$	$\hat{\beta}_{\text{S/GLILr}}^{\text{thr}}$
Stage 1	...	8.5	−1
Stage 2	...	10	−1
Stage 3	...	14	−5
Stage 4	3	...	...
Stage 5	7	...	...
Stage 6	15	...	...
Stage 7	35	...	...
Stage 8	15	...	...

amplitude h_0 at the detector. This is the smallest continuous gravitational-wave amplitude such that we can recover 90% of the signals of our target population. We estimate the upper limits in half-Hz bands with injection-and-recovery Monte Carlo simulations and show the results in Figure 8. These are available as the data behind Figure 8.

We employ the same method as in Steltner et al. (2021). In a first pass, we add fake signals drawn from the target population, treat the resulting data exactly as is done in the data preparation for the search, search with the Stage 0 setup, and cluster the results as done for the Stage 0 results. In a second pass, we do not add the fake signal, and perform every step exactly alike. We consider the signal to be detected if the results with the added fake signal (i) produce a candidate with parameters close enough (within the uncertainty region) to the signal parameters, and (ii) the detection statistic value of that candidate is larger than the detection statistic measured at the same parameter space point in the absence of a signal. We do this in every half-Hz band 200 times for each h_0 value. The confidence $C(h_0)$ is the fraction of detected signals at that h_0 value. We consider at least five different h_0 values. We fit the resulting $C(h_0)$ curve and derive the upper limit value h_0^{UL} such that $C(h_0^{\text{UL}}) = 0.9$, as described in Section 5.1 of Fesik & Papa (2020).




Figure 7. The results of each follow-up stage. The top three right-hand-side panels show a scatter plot of the $(2\bar{\mathcal{F}}_r, \hat{\beta}_{\text{S/GLILr}})$ values of candidates from the search and candidates from Monte Carlo simulations containing fake signals. The region below the black lines is the candidate rejection region. The wide panels show the distribution of R^a for search candidates and for candidates from Monte Carlo simulations containing fake signals. The dashed vertical lines indicate the value of $R^{a-\text{thr}}$. The dashed horizontal line marks the level where one candidate appears in the normalized histogram.

Our upper limits do not hold in the 50 mHz bands that we marked as disturbed. As already said, the disturbed band list is provided in a supplemental .tar.gz package. Upper limits are also not given in 0.5 Hz bands where (i) all 50 mHz bands are marked disturbed or (ii) where the 90% detection efficiency is not reached, due to the line-cleaning procedure removing the added fake signal, as it would have done with a real signal. We find that the upper limit value is impacted by the cleaning in less than 6% of the half-Hz bands.

Based on our upper limits, we achieve sensitivity depths in the range of $(52 - 59)1/\sqrt{\text{Hz}}$. This constitutes a sensitivity improvement of $\sim 6\%$ stemming solely from the search method, with respect to the previous O2 Einstein@Home search (Steltner et al. 2021).

The upper limits on h_0 can be translated into upper limits on the ellipticity ε of a source modeled as a triaxial ellipsoid spinning around a principal moment of inertia axis I at a distance d :

$$\varepsilon^{\text{UL}} = 1.4 \times 10^{-6} \left(\frac{h_0^{\text{UL}}}{1.4 \times 10^{-25}} \right) \times \left(\frac{d}{1 \text{ kpc}} \right) \left(\frac{170 \text{ Hz}}{f} \right)^2 \left(\frac{10^{38} \text{ kg m}^2}{I} \right). \quad (9)$$

Table 3

Frequency, Frequency Derivative, and Sky Position of the Hardware Injections and the Distance between These and the Values Recovered by Our Coherent Search Using the Entire O3 Data (Stage 9)

ID_{inj}	f (Hz)	$\dot{f}$ (Hz s ⁻¹)	α (hr:m:s)	δ (deg:m:s)	Δf (Hz)	$\Delta \dot{f}$ (Hz s ⁻¹)	Sky Distance (deg:m:s)
0	265.57505348	-4.15×10^{-12}	4:46:12.4628	-57:46:57.0510	-4.7×10^{-11}	9.5×10^{-16}	0:0:0.0741
2	575.16350527	-1.37×10^{-13}	14:21:1.4800	3:26:38.3626	-1.1×10^{-9}	-8.8×10^{-16}	0:0:0.0955
3	108.85715939	-1.46×10^{-17}	11:53:29.4178	-34:33:48.2313	6.7×10^{-10}	-5.8×10^{-16}	0:0:0.3080
5	52.80832436	-4.03×10^{-18}	20:10:30.3939	-84:9:39.0964	-6.2×10^{-10}	-4.2×10^{-16}	0:0:0.2212
9	763.84731649	-1.45×10^{-17}	13:15:32.5397	75:41:22.5205	9.4×10^{-10}	-5.6×10^{-17}	0:0:0.0023
10	26.33209638	-8.50×10^{-11}	14:46:13.3549	42:52:38.2953	-8.3×10^{-11}	2.4×10^{-16}	0:0:0.3109

Note. The frequency f is given at the reference time 1253764756.0 (GPS time).




Figure 8. The smallest gravitational-wave amplitude h_0 that we can exclude from our target population. We compare to other all-sky searches in LIGO O3 data (Abbott et al. 2022; Dergachev & Papa 2022). There are multiple curves for Abbott et al. (2022), corresponding to multiple analysis pipelines. The golden stars are the gravitational-wave amplitudes h_0 of the hardware injections. Of the seven hardware injections in our search range, we recover all but hardware injection 11 (red star), due to its parameters being at a high mismatch point in the search grid. To miss one injection of seven is perfectly consistent with our 90% confidence upper limits. The data from this work are available as the data behind the figure.

(The data used to create this figure are available.)

Figure 9 shows the upper limits on the ellipticity ε for different distances.

Another possible emission mechanism is r -modes, unstable toroidal fluid oscillations driven by the Coriolis force, emitting at $\approx 4/3$ of the spin frequency (Andersson 1998; Friedman & Morsink 1998; Lindblom et al. 1998). We translate the upper limits on the gravitational-wave amplitude h_0 into upper limits on the r -mode amplitude (Owen 2010):

$$\alpha^{\text{UL}} = 0.028 \left(\frac{h_0^{\text{UL}}}{1 \times 10^{-24}} \right) \left(\frac{d}{1 \text{ kpc}} \right) \left(\frac{100 \text{ Hz}}{f} \right)^3. \quad (10)$$

Figure 10 shows the upper limits on the r -mode amplitude α for different distances d .




Figure 9. Upper limits on the neutron star ellipticity ε at different distances. The dashed line shows the maximum ellipticity probed due to the maximum spin-down of this search. The slightly increased upper limit values near 120 Hz, 305 Hz, and 435 Hz—also visible in the h_0 upper limit plot—are due to decreased search sensitivity, due to the line-cleaning procedure.

7. Conclusions

We present the results from an Einstein@Home all-sky search for continuous gravitational waves with frequency between $20.0 \text{ Hz} \leq f \leq 800.0 \text{ Hz}$ and a spin-down of $-2.6 \times 10^{-9} \text{ Hz s}^{-1} \leq \dot{f} \leq 2.6 \times 10^{-10} \text{ Hz s}^{-1}$ in the LIGO O3 public data.

Of the 6.7×10^{18} waveforms searched, we identify the most promising 3.5 million and investigate them with a nine-stage hierarchical follow-up. The only surviving candidates are associated to the hardware injection fake signals.

We recover all hardware injections that lie above/at our upper limit level and even two with amplitude significantly below the upper limit curve. All signal parameters are recovered with high accuracy.

The most stringent 90% confidence level upper limit on the gravitational-wave amplitude h_0 is placed at 203 Hz, at the level of 8.1×10^{-26} . This excludes neutron stars with ellipticity ε , at a distance d and rotating with periods P :

$$P < 5 \text{ ms} \quad \text{and} \quad \varepsilon > 5 \times 10^{-8} \left[\frac{d}{100 \text{ pc}} \right]. \quad (11)$$

Depending on the assumptions, the nearest unknown neutron star could be as close as $\approx 10 \text{ pc}$ (Dergachev & Papa 2020; Pagliaro et al. 2023), where this search probes the interesting region of ellipticities between 10^{-7} and 10^{-9} for sufficiently




Figure 10. Upper limits on the r -mode amplitude α at different distances.

fast-spinning (≥ 50 Hz) objects. Pagliaro et al. (2023), however, find that the average distance of the nearest neutron star *spinning in band* is approximately 100 pc. This means that, in practice, we are most likely probing ellipticities between 10^{-6} and 10^{-8} .

It is hence natural to ask if objects with ellipticities larger than 10^{-8} could exist. There is much uncertainty about the maximum ellipticity that a neutron star crust can support, and the mechanisms to create such ellipticity, especially for isolated objects. While some models predict maximum ellipticities around $\sim 10^{-5}$ (Johnson-McDaniel & Owen 2013; Morales & Horowitz 2022), other works show that the maximum may be only at 10^{-9} (Bhattacharyya 2020; Gittins & Andersson 2021). Overall, ellipticities at the level of 10^{-6} are usually considered reasonable, and the reach of this search for objects with this deformation and rotating faster than 150 Hz is 1 kpc.

Young, energetic neutron stars are promising emitters of gravitational waves due to r -modes. The amplitudes probed here are not implausible for a star a few years—at most, a few decades—old, born anywhere in the Galaxy (Bondaescu et al. 2009).

This is the most sensitive all-sky search performed on this parameter space to date. Still, we cannot claim that with the same computational budget, an even more sensitive search could not be performed. Rather, our search setup investigations did not identify one. Our investigations are resource-intensive, because they rely on Monte Carlo simulations aimed at *measuring* the detection efficiency of signals under different search conditions. And because they are resource-intensive, the scope of the considered setups is limited. We would like to overcome these limitations in the future. One approach is to make the simulations more efficient; alternatively, one could use an analytical model to predict the sensitivity of the search. Such a model would have to be accurate to a few percent, which is the scale of the sensitivity depth improvements among different methods working on the same data. Given the nonlinearity of many of the steps (clustering, to mention one), obtaining this accuracy in the prediction of the sensitivity is not trivial.

However the search setup will be arrived at, we look forward to the LIGO/Virgo data of the next observing run (O4), which

we will use to probe an even more interesting range of deformations/distances.

Acknowledgments

We thank the Einstein@Home volunteers, without the support of whom this search could not have happened. We thank Gianluca Pagliaro for useful discussions on the detectability of Galactic neutron stars. We acknowledge the use of Topcat (Taylor 2005). B.A and M.A.P. acknowledge NSF award grant No. 1816904 for support of the Einstein@Home project. This research has made use of data or software obtained from the Gravitational Wave Open Science Center (gwosc.org), a service of LIGO Laboratory, the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA.

ORCID iDs

B. Steltner <https://orcid.org/0000-0003-1833-5493>
M. A. Papa <https://orcid.org/0000-0002-1007-5298>
H.-B. Eggenstein <https://orcid.org/0000-0001-5296-7035>
B. Allen <https://orcid.org/0000-0003-4285-6256>

References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, *PhRvD*, **96**, 122004
- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019, *PhRvX*, **9**, 031040
- Abbott, R., Abbott, T. D., Abraham, S., et al. 2021a, *PhRvX*, **11**, 021053
- Abbott, R., Abbott, T. D., Abraham, S., et al. 2021c, *SoftX*, **13**, 100658
- Abbott, R., Abe, H., Acernese, F., et al. 2022, *PhRvD*, **106**, 102008
- Acernese, F., Agathos, M., Agatsuma, K., et al. 2015, *CQGra*, **32**, 024001
- Anderson, D. P. 2004, in Proc. Fifth IEEE/ACM Int. Workshop on Grid Computing (GRID04) (Washington, DC: IEEE Computer Society), 4
- Anderson, D. P., Christensen, C., & Allen, B. 2006, in Proc. 2006 ACM/IEEE Conf. on Supercomputing, ed. B. Horner-Miller (New York: ACM), 126
- Andersson, N. 1998, *ApJ*, **502**, 708
- Arvanitaki, A., Baryakhtar, M., & Huang, X. 2015, *PhRvD*, **91**, 084011
- Behnke, B., Papa, M. A., & Prix, R. 2015, *PhRvD*, **91**, 064007
- Bhattacharyya, S. 2020, *MNRAS*, **498**, 728
- Bondaescu, R., Teukolsky, S. A., & Wasserman, I. 2009, *PhRvD*, **79**, 104003
- Brady, P. R., & Creighton, T. 2000, *PhRvD*, **61**, 082001
- Cutler, C., & Schutz, B. F. 2005, *PhRvD*, **72**, 063006
- Davis, D., Massinger, T. J., Lundgren, A. P., et al. 2019, *CQGra*, **36**, 055011
- Dergachev, V., & Papa, M. A. 2020, *PRL*, **125**, 171101
- Dergachev, V., & Papa, M. A. 2022, arXiv:2202.10598
- Fesik, L., & Papa, M. A. 2020, *ApJ*, **895**, 11
- Friedman, J. L., & Morsink, S. M. 1998, *ApJ*, **502**, 714
- Gittins, F., & Andersson, N. 2021, *MNRAS*, **507**, 116
- Horowitz, C., Papa, M., & Reddy, S. 2020, *PhLB*, **800**, 135072
- Horowitz, C. J., & Reddy, S. 2019, *PhRvL*, **122**, 071102
- Jaranowski, P., Krolak, A., & Schutz, B. F. 1998, *PhRvD*, **58**, 063001
- Johnson-McDaniel, N. K., & Owen, B. J. 2013, *PhRvD*, **88**, 044004
- Keitel, D. 2016, *PhRvD*, **93**, 084024
- Lasky, P. D. 2015, *PASA*, **32**, e034
- LIGO, S. C., & Virgo, P. 2022, https://www.gw-openscience.org/O3/O3April1_injection_parameters/
- Lindblom, L., Owen, B. J., & Morsink, S. M. 1998, *PhRvL*, **80**, 4843
- Morales, J. A., & Horowitz, C. J. 2022, *MNRAS*, **517**, 5610
- Nitz, A. H., Capano, C., Nielsen, A. B., et al. 2019, *ApJ*, **872**, 195
- Nitz, A. H., Capano, C. D., Kumar, S., et al. 2021a, *ApJ*, **922**, 76
- Nitz, A. H., Dent, T., Davies, G. S., et al. 2020, *ApJ*, **891**, 123
- Nitz, A. H., Kumar, S., Wang, Y.-F., et al. 2021b, arXiv:2112.06878
- Olsen, S., Venumadhav, T., Mushkin, J., et al. 2022, *PhRvD*, **106**, 043009
- Owen, B. J. 2010, *PhRvD*, **82**, 104002
- Owen, B. J., Lindblom, L., Cutler, C., et al. 1998, *PhRvD*, **58**, 084020
- Pagliaro, G., Papa, M. A., Ming, J., et al. 2023, arXiv:2303.04714
- Pletsch, H. J. 2010, *PhRvD*, **82**, 042002
- Singh, A., Papa, M. A., Eggenstein, H.-B., & Walsh, S. 2017, *PhRvD*, **96**, 082003

- Steltner, B., Menne, T., Papa, M. A., & Eggenstein, H.-B. 2022a, [PhRvD](#), **106**, 104063
- Steltner, B., Papa, M. A., Eggenstein, H. B., et al. 2021, [ApJ](#), **909**, 79
- Steltner, B., Papa, M. A., & Eggenstein, H.-B. 2022b, [PhRvD](#), **105**, 022005
- Taylor, M. B. 2005, in ASP Conf. Ser. 347, Astronomical Data Analysis Software and Systems XIV, ed. P. Shopbell, M. Britton, & R. Ebert (San Francisco, CA: ASP), 29
- The LIGO Scientific Collaboration, Aasi, J., Abbott, B. P., et al. 2015, [CQGra](#), **32**, 074001
- The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration 2023, arXiv:2302.03676
- The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration 2021b, arXiv:2111.03606
- Vajente, G., Huang, Y., Isi, M., et al. 2020, [PhRvD](#), **101**, 042003
- Venumadhav, T., Zackay, B., Roulet, J., Dai, L., & Zaldarriaga, M. 2019, [PhRvD](#), **100**, 023011
- Venumadhav, T., Zackay, B., Roulet, J., Dai, L., & Zaldarriaga, M. 2020, [PhRvD](#), **101**, 083030
- Zhu, S. J., Baryakhtar, M., Papa, M. A., et al. 2020, [PRD](#), **102**, 063020