

Constraints on gravitational waves from the 2024 Vela pulsar glitch

THE LIGO SCIENTIFIC COLLABORATION, THE VIRGO COLLABORATION, THE KAGRA COLLABORATION, JIM PALFREYMAN ,¹ S.B. ARAUJO FURLAN ,^{2,3} S. DEL PALACIO ,^{4,5} G. GANCIO ,⁴ F. GARCÍA ,^{4,6} G. E. ROMERO ,^{4,6} AND E. ZUBIETA ,^{4,6}

¹*School of Natural Sciences, University of Tasmania, Hobart, Australia*

²*Instituto de Astronomía Teórica y Experimental, CONICET-UNC, Laprida 854, X5000BGR – Córdoba, Argentina.*

³*Facultad de Matemática, Astronomía, Física y Computación, UNC. Av. Medina Allende s/n , Ciudad Universitaria, CP:X5000HUA - Córdoba, Argentina.*

⁴*Instituto Argentino de Radioastronomía (CCT La Plata, CONICET; CICPBA; UNLP), C.C.5, (1894) Villa Elisa, Buenos Aires, Argentina*

⁵*Department of Space, Earth and Environment, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden*

⁶*Facultad de Ciencias Astronómicas y Geofísicas, Universidad Nacional de La Plata, Paseo del Bosque, B1900FWA La Plata, Argentina*

ABSTRACT

Among known neutron stars, the Vela pulsar is one of the best targets for gravitational-wave searches. It is also one of the most prolific in terms of glitches, sudden frequency changes in a pulsar’s rotation. Such glitches could cause a variety of transient gravitational-wave signals. Here we search for signals associated with a Vela glitch on 29 April 2024 in data of the two LIGO detectors from the fourth LIGO–Virgo–KAGRA observing run. We search both for seconds-scale burst-like emission, primarily from fundamental (f-)mode oscillations, and for longer quasi-monochromatic transients up to four months in duration, primarily from quasi-static quadrupolar deformations. We find no significant detection candidates, but for the first time we set direct observational upper limits on gravitational strain amplitude that are stricter than what can be indirectly inferred from the overall glitch energy scale. We discuss the short- and long-duration observational constraints in the context of specific emission models. These results demonstrate the potential of gravitational-wave probes of glitching pulsars as detector sensitivity continues to improve.

Draft revision: git hash 78f70416 (December 19, 2025)

Keywords: Gravitational wave astronomy(675) — Neutron stars(1108) — Pulsars(1306) — Interdisciplinary astronomy(804)

1. INTRODUCTION

Pulsar glitches are sudden, enigmatic jumps in a pulsar’s spin frequency, which have been observed so far in over 200 pulsars (Espinoza et al. 2011; Yu et al. 2013; Basu et al. 2022). They generally involve an initial rapid change in frequency and frequency derivative (on timescales of less than a minute) followed by a more gradual (on timescales of up to months) relaxation towards the previous spin-down state (Antonopoulou et al. 2022). While the common view is that glitches are linked to an exchange of angular momentum between an interior superfluid component and the crust of a neutron star (NS), the quantitative details are still an open question. Diverse mechanisms such as vortex avalanches, fluid instabilities and crust-quakes have been invoked as possible triggers (see Haskell & Melatos 2015; Zhou et al. 2022, for reviews).

Glitches are key tools for studying NS interiors (Link et al. 1992, 2000; Haskell 2018) and could also trigger various types of gravitational-wave (GW) emission: both transient signals linked to the initial jump, and longer-lived signals

linked to the relaxation (Andersson & Kokkotas 1998; van Eysden & Melatos 2008; Bennett et al. 2010; Prix et al. 2011; Keer & Jones 2015; Melatos et al. 2015; Singh 2017; Ho et al. 2020; Cheunchitra et al. 2024; Yim & Jones 2020, 2023; Wilson & Ho 2024; Yim et al. 2024; Haskell & Jones 2024). Studying these signals can constrain the NS mass and radius and hence the nuclear matter equation of state.

Among the known pulsar population, the Vela pulsar has some of the largest and most frequent glitches (Fuentes et al. 2017). It is also relatively close to Earth at a distance of 287_{-17}^{+19} pc (Dodson et al. 2003). For these reasons, it has been identified in several studies (Moragues et al. 2023; Lopez et al. 2022; Yim et al. 2024) as a priority target to search for glitch-associated transient GWs.

Here, we present a set of searches for GW emission associated with the 29 April 2024 glitch of the Vela pulsar, first observed (Zubieta et al. 2024a, 2025) by the Argentine Institute of Radio Astronomy (IAR) (Zubieta et al. 2023) and subsequently confirmed by several other observatories and precisely timed (Palfreyman 2024) by the Mount Pleas-

ant Radio Observatory (MPRO) of the University of Tasmania. This glitch happened during the LIGO–Virgo–KAGRA (LVK) fourth observing run (O4), shortly after the two LIGO detectors (Aasi et al. 2015) had restarted observations for the second part (O4b) of the run and were joined also by Virgo (Acernese et al. 2015). The electromagnetic (EM) observations of the glitch are summarized in Section 2 and the GW data set is described in Section 3.

To cover a range of glitch-related GW emission scenarios (see Haskell & Jones 2024 for a review), our searches fall into two categories: First, we search for transient GW bursts that may be triggered by the initial rapid rise in frequency, on dynamical timescales of milliseconds to seconds, primarily from f-mode oscillations at kHz frequencies, in Section 4. Then, we cover longer-duration (up to four months) quasi-monochromatic signals at lower frequencies, coupled with the pulsar spin frequency and associated with the post-glitch relaxation, in Section 5. Even longer signals are well covered by standard continuous wave (CW) searches. Details of the emission models are included in each of these sections.

Previous searches for GWs from pulsar glitches have included an f-mode burst search on the August 2006 Vela glitch using initial LIGO data (Abadie et al. 2011), various all-sky burst searches (latest O4a results: Abac et al. 2025a,b), and long-duration searches on O2 data (Keitel et al. 2019; Modafferi et al. 2023), including the 2016 Vela glitch (Palfreyman et al. 2018), and O3 data (Abbott et al. 2022a; Modafferi et al. 2021, not including a Vela glitch). However, matching the studies of detection prospects for the two cases by Lopez et al. (2022); Moragues et al. (2023), none of these searches had yet been able to make a GW detection or to place physically constraining observational upper limits below the energy scale indirectly set by the glitch size. Searches for NS f-modes have also been performed on bursting magnetars, with no detection yet either (e.g. Abbott et al. 2024).

The improved sensitivity of the LIGO O4 data (Capote et al. 2025; Ganapathy et al. 2023; Jia et al. 2024) now gives us for the first time the opportunity to reach physically constraining results on GWs from a pulsar glitch. We discuss their implications for pulsar glitch models and GW emission channels in Section 6 and present our conclusions in Section 7.

2. THE VELA PULSAR AND ITS 2024 GLITCH

The Vela Pulsar (PSR B0833–45/J0835–4510), discovered by Large et al. (1968), was the first pulsar known to glitch (Radhakrishnan & Manchester 1969). It is young ($\sim 10^4$ yr) and one of the most active pulsars in terms of glitching, with 23 large glitches (often considered as those with relative frequency changes $\Delta f_{\text{rot}}/f_{\text{rot}} > 10^{-7}$, Espinoza et al. 2011; Yu et al. 2013; Basu et al. 2022) reported so far – approximately once every two to three years (Re-

ichley & Downs 1969; Dodson et al. 2007). It also exhibits micro-glitches sporadically (Cordes et al. 1988; Espinoza et al. 2021; Dunn et al. 2023; Zubieta et al. 2024b).

During O4, the Vela pulsar has been monitored with almost daily combined cadence by the IAR and MPRO.

2.1. IAR observations

The IAR¹ was founded in 1962 as a pioneer radio observatory in South America. Located near the city of La Plata, Buenos Aires, it is located at $34^{\circ}51'57''.35$ S $58^{\circ}08'25''.04$ W. It has two 30 m diameter single-dish antennas, aligned on a North–South direction, separated by 120 m. Observing around 1400 MHz with 56 MHz and 400 MHz bandwidths simultaneously, the data is obtained with a time resolution of 146 μ s.

A summary of IAR pulsar observation work and future projects is given in Gancio et al. (2020); Lousto et al. (2024) and references therein. The IAR has reported the observation of the last three large Vela glitches: on 1 February 2019 with a relative frequency change of 2.7×10^{-6} (Lopez Armengol et al. 2019; Gancio et al. 2020), on 22 July 2021 with relative change of 1.2×10^{-6} (Sosa-Fiscella et al. 2021), and the 29 April 2024 glitch of interest for this paper (counted as large Vela glitch #23) first reported in Zubieta et al. (2024a) with $(\Delta f_{\text{rot}}/f_{\text{rot}}) = 2.4 \times 10^{-6}$ and a more detailed analysis given in Zubieta et al. (2025).

For the data used in this paper, we used the PRESTO (Ransom 2011) package to remove radio-frequency interferences (RFIs) and fold the observations. We cross-correlated the folded profile with a noise-free template profile in order to obtain the Time of Arrivals (TOAs). We obtained the template with the tool `psrsmooth` in PSRCHIVE (Hotan et al. 2004) on a pulse profile with high signal-to-noise ratio (S/N).

2.2. MPRO observations

The University of Tasmania’s MPRO 26 m radio telescope is located at Mt. Pleasant, near Hobart, Australia at $42^{\circ}48'12''.90$ S $147^{\circ}26'25''.87$ E. It has previously detected several Vela glitches, including the real-time observation of the 2016 glitch with associated magnetospheric changes (Palfreyman et al. 2018), which suggest a connection between the magnetosphere and the internal dynamics of the pulsar that probably provokes glitches.

The center frequency of the observations used here is 2250 MHz with bandwidth 64 MHz. Processing was done using the software packages DSPSR (van Straten & Bailes 2011), TEMPO2 (Hobbs et al. 2006; Edwards et al. 2006), and PSRCHIVE, which perform phase-coherent de-dispersion using 16 frequency channels and 8192 pulse phase

¹ <https://www.iar.unlp.edu.ar>

bins over the 89.3 ms pulse period. This gives a time resolution of 10.9 μ s.

2.3. Timing model

We collected the TOAs from both observatories between Modified Julian Date (MJD) 60402.28 and MJD 60493.53 and characterized the rotation of the Vela pulsar with TEMPO2 by fitting a Taylor series

$$\phi(t) = f_{\text{rot}}(t - T_{\text{ref}}) + \frac{1}{2} \dot{f}_{\text{rot}}(t - T_{\text{ref}})^2 + \dots \quad (1)$$

The pulsar's rotation frequency is f_{rot} , its derivatives are \dot{f}_{rot} , \ddot{f}_{rot} , etc., and T_{ref} is the reference epoch.

To characterize the glitch, we added a glitch-induced change in the rotational phase (Edwards et al. 2006; Yu et al. 2013) to this model:

$$\begin{aligned} \phi_g(t) = & \Delta\phi + \Delta f_{\text{rot}}(t - T_{\text{gl}}) + \frac{1}{2} \Delta \dot{f}_{\text{rot}}(t - T_{\text{gl}})^2 \\ & + \sum_i \left[1 - \exp\left(-\frac{t - T_{\text{gl}}}{\tau_d^i}\right) \right] \Delta f_{\text{rot}}^{d,i} \tau_d^i. \end{aligned} \quad (2)$$

Here, T_{gl} is the epoch of the glitch, with its uncertainty counteracted by $\Delta\phi$, while Δf_{rot} and $\Delta \dot{f}_{\text{rot}}$ are the sudden changes in f_{rot} and \dot{f}_{rot} at T_{gl} . In addition, $\Delta f_{\text{rot}}^{d,i}$ are the temporary increases (recovery terms) in the rotational frequency of the pulsar that decay exponentially on a timescale τ_d^i days.

The 29 April 2024 Vela glitch is evident from the pre-fit residuals in Figure 1. We characterized the pre-glitch timing model by fitting Equation (1) to TOAs from one month of observations before the glitch. By combining the IAR and MPRO data, we found preference for three transient increase terms in the frequency of the pulsar in addition to the permanent changes Δf_{rot} and $\Delta \dot{f}_{\text{rot}}$. The post-fit residuals can be seen in the bottom panel of Figure 1. We then extended the data span to MJD 60357–60551, which is the full range of interest for the GW search here. With this data span, we fitted f_{rot} , \dot{f}_{rot} , \ddot{f}_{rot} and $\ddot{\dot{f}}_{\text{rot}}$ keeping the glitch parameters frozen. The full timing solution covering the extended range of observations is shown in Table 1. This is valid within the listed START and FINISH dates, and hence the listed \ddot{f}_{rot} should not be interpreted as giving the long term braking index (see Vargas & Melatos 2023).

3. GW DATA

The O4 run of the advanced laser-interferometric GW detectors began on 27 May 2023 and, after a commissioning break with several interventions to further improve sensitivity, its second part O4b ran from 10 April 2024 to 28 January 2025. In O4b, the two LIGO detectors (Aasi et al. 2015) in Hanford (Washington, US) and Livingston (Louisiana, US)

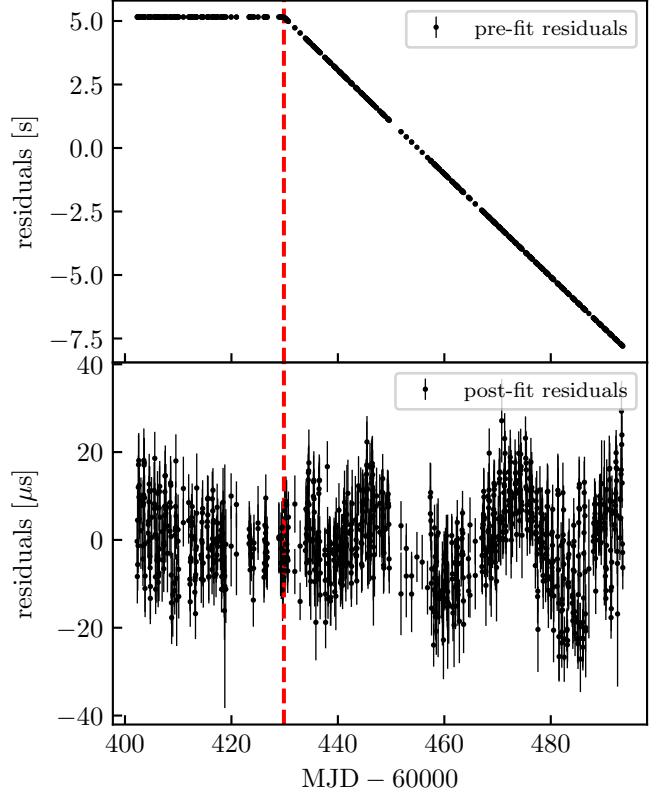


Figure 1. Top panel: Vela pulsar radio timing residuals before fitting for the 29 April 2024 glitch. Bottom panel: Timing residuals after fitting for the glitch with the parameters shown in Table 1.

and the Virgo detector (Acernese et al. 2015) in Cascina (Italy) participated. KAGRA (Akutsu et al. 2021) in Japan rejoined O4 later. The upgrades to LIGO for O4 allowed it to reach unprecedented GW strain sensitivity (Capote et al. 2025; Ganapathy et al. 2023; Jia et al. 2024).

At the time of the Vela glitch on 29 April 2024, the LIGO Hanford (H1), LIGO Livingston (L1) and Virgo (V1) detectors were all operating in a stable state; see Figure 2. The burst-type searches in Section 4 use data within two days of the glitch time. The coincident data taking of the two LIGO detectors with similar sensitivity across most of the frequency band was essential, as a single-detector search of this type loses an order of magnitude in sensitivity (e.g., Ball & Frey 2025) due to missing suppression capability towards instrumental glitches (not to be confused with pulsar glitches).

The longer-duration searches in Section 5 cover a maximum extent starting one day before the glitch epoch and ending on 28 August 2024. For part of this time stretch, only L1 data are available, as H1 had to be taken offline to diagnose and repair a problem in the output optics chain from 13 July to 24 August, and V1 went into a downtime for limiting noise factors from 22 July to 2 August. The analyzable LIGO data

Table 1. Parameters of the Vela pulsar and its 2024 glitch.

parameter	value
R.A.	08:35:20.61993 \pm 0.00002
DEC.	-45:10:33.7232 \pm 0.0003
d [pc]	287 $^{+19}_{-17}$ ^(a)
T_{ref} [d] (MJD)	60408
T_{ref} [s] (GPS)	1396569549
START (MJD)	60357
FINISH (MJD)	60551
f_{rot} [Hz]	(11.182888898688 \pm 0.000000000002)
f_{rot} [Hz s $^{-1}$]	(-1.5541536 \pm 0.000001) \times 10 $^{-11}$
\ddot{f}_{rot} [Hz s $^{-2}$]	(3.82 \pm 0.03) \times 10 $^{-22}$
$\ddot{\ddot{f}}_{\text{rot}}$ [Hz s $^{-3}$]	(1.283 \pm 0.07) \times 10 $^{-28}$
T_{gl} [d] (MJD)	60429.86975 \pm 0.00139
T_{gl} [s] (GPS)	1398459095.216 \pm 120.096
$\Delta\phi$	-0.00034 \pm 0.00017
Δf_{rot} [Hz]	(2.65854 \pm 0.00005) \times 10 $^{-5}$
$\Delta \dot{f}_{\text{rot}}$ [Hz s $^{-1}$]	(-1.0298 \pm 0.0009) \times 10 $^{-13}$
$\Delta \ddot{f}_{\text{rot}}$ [Hz s $^{-2}$]	—
τ_d^1 [d]	15.1 \pm 0.1
$\Delta f_{\text{rot}}^{d,1}$ [Hz]	(1.501 \pm 0.004) \times 10 $^{-7}$
τ_d^2 [d]	2.45 \pm 0.06
$\Delta f_{\text{rot}}^{d,2}$ [Hz]	(1.24 \pm 0.03) \times 10 $^{-7}$
τ_d^3 [d]	0.39 \pm 0.03
$\Delta f_{\text{rot}}^{d,3}$ [Hz]	(1.8 \pm 0.1) \times 10 $^{-7}$

NOTE— (a) [Dodson et al. \(2003\)](#)

All other values were derived from IAR and MPRO observations, as described in Section 2.

Reference time and glitch epoch (at the solar system barycenter) are given in both MJD (terrestrial time standard), as is conventional in radio astronomy, and in GPS seconds (UTC standard) for LVK data analysis. T_{gl} is given to the precision ingested by burst analyses.

segments ([Goetz & Riles 2024](#)) over this extended timescale are also shown in Figure 2.

Due to the large difference in sensitivity, in terms of the detectors’ ASDs, Virgo data is not used in this paper. From LIGO, we use data from the GDS-CALIB_STRAIN_CLEAN frame channel, corresponding to the standard online calibration ([Viets et al. 2018; Sun et al. 2020; Sun et al. 2021; Wade et al. 2025](#)) with some noise subtraction applied ([Viets & Wade 2021; Vajente et al. 2020; Abac et al. 2025c](#)). For several searches in Section 5, additional time-domain removal of large noise excursions was performed ([Astone et al. 2005; Davis et al. 2025](#)), before creating Fourier-domain data products as described later for each analysis method. Data quality

has been investigated with methods similar to those described for O4a in [Soni et al. \(2025\)](#).

4. SEARCHES FOR GW BURSTS

In the context of this paper, we refer to any signals as “GW bursts” that last from milliseconds up to, at most, a few minutes. The main physical emission channel considered are NS f-modes excited by the glitch and damping away again on the signal timescale ([Andersson & Kokkotas 1998](#)). However, all three search methods that we use here are unmodeled or very weakly modeled, i.e., they do not rely on matched filtering. Thus, they are also sensitive to other emission mechanisms in the same time and frequency ranges. The allowed signal durations and the on-source windows (time range in which the signal might start) for the three methods are summarized in Figure 3. Here, the glitch time is the GPS time at the detector.

These, or similar, methods have been routinely used for all-sky GW burst searches in LVK data (e.g., [Abbott et al. 2021a,b; Abac et al. 2025a,b](#)), which would have also been sensitive to signals from pulsar glitches with sufficiently high amplitudes ([Lopez et al. 2022](#)). However, the last targeted search for GW bursts from a pulsar glitch dates back to the era of the initial LIGO detectors, when the 2006 Vela glitch was analyzed ([Abadie et al. 2011](#)) using the two Hanford detectors available at the time. With the two Advanced LIGO detectors and improved analysis methods, we now obtain orders of magnitude better results than that study.

In this section, after briefly discussing the f-mode scenario, we introduce the three pipelines used and summarize their results. They cover different time and frequency ranges to increase robustness to different emission scenarios. For each search, after finding no significant candidates, we conduct injections of simulated signals into the GW data to estimate search sensitivity and provide upper limits; this is discussed in Section 4.5. Detailed interpretation is left to Section 6.

4.1. Emission Models

Glitching pulsars can emit short-duration GW signals via mechanisms such as an avalanche of unpinning superfluid vortices (see, e.g., [Warszawski & Melatos 2012; De Lillo et al. 2023](#), and references therein) or oscillation modes of the NS ([Andersson & Kokkotas 1998; Keer & Jones 2015; Yim & Jones 2023; Wilson & Ho 2024](#)).

The situation is complicated, as NSs can support a variety of oscillations modes, classified according to the dominant restoring force, e.g. pressure, buoyancy, or rotation ([Kokkotas & Schmidt 1999](#)). We mostly focus on the fundamental (f-)modes, which are good candidates for our searches for several reasons: They are well understood, with frequencies depending only on the gross features of the NS equation of state, typically in the kHz range accessible by LVK detectors. They are also extremely efficient GW emitters, in the sense

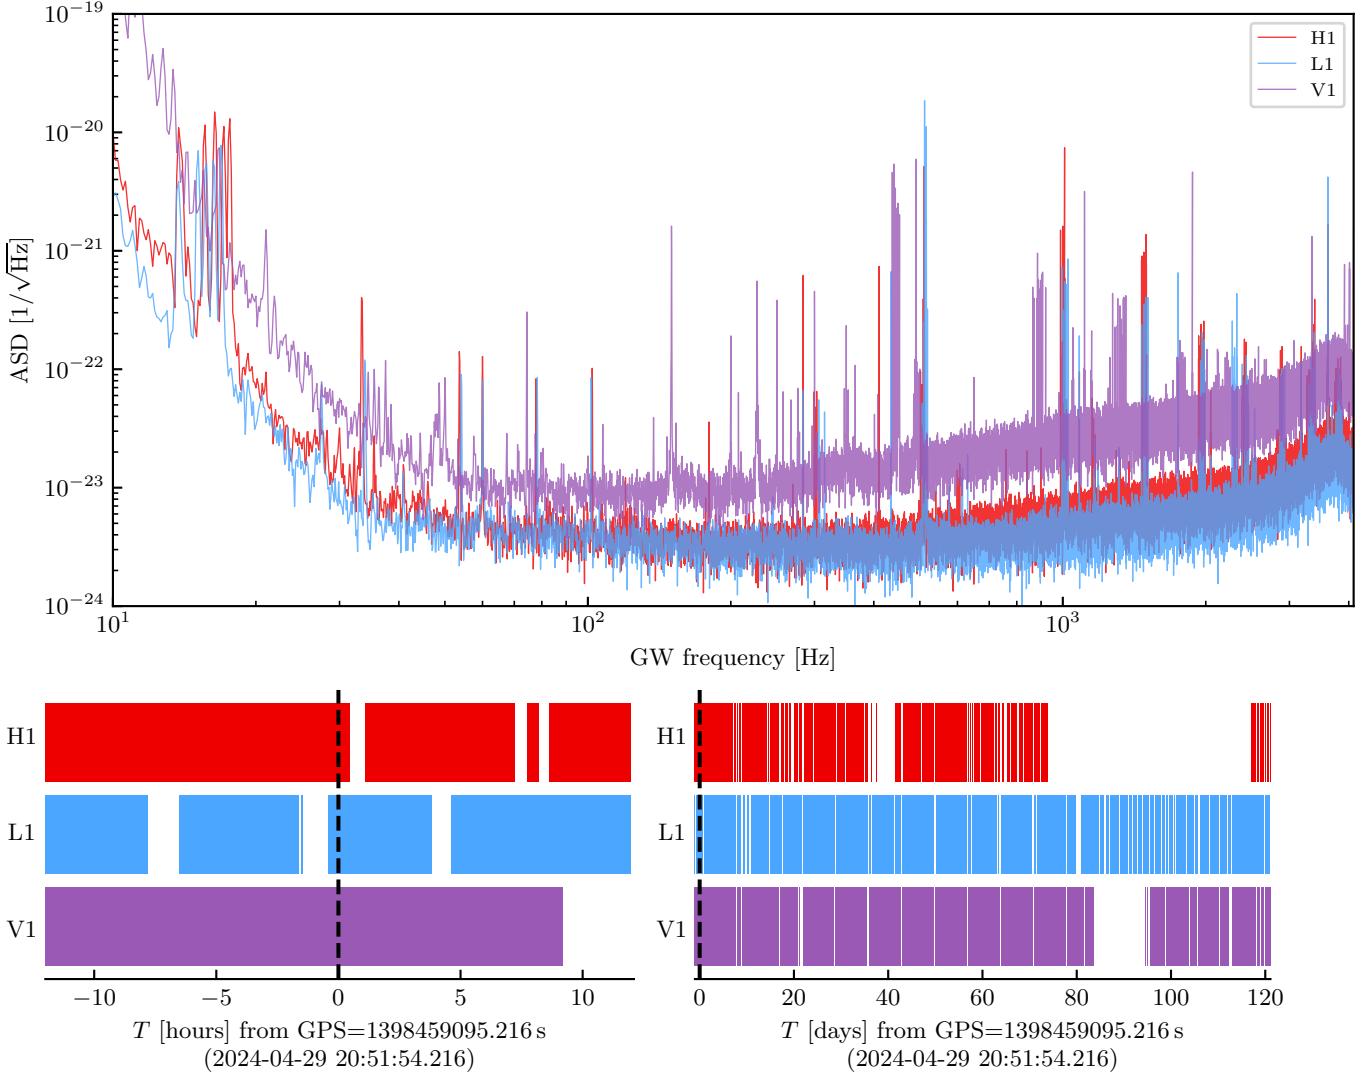


Figure 2. Top panel: GW detector sensitivities within $T_{\text{gl}} \pm 20$ s around the Vela glitch on 29 April 2024, in terms of their ASDs (computed with the Welch method in `gwpy`, Macleod et al. 2021). Bottom left panel: GW data availability over 24 h around the glitch, as used for the burst searches in Section 4. Bottom right panel: LIGO data availability for the four following months, as used in the long-duration searches in Section 5. Virgo data is not used in either type of searches. The glitch time is highlighted in the two bottom panels with a vertical dashed line.

that almost all energy deposited in the mode by the glitch is emitted in quadrupolar GWs, with very little energy loss to viscosity (Echeverria 1989; Finn 1992; Ho et al. 2020).

Regardless of the precise emission mechanism, we can define a *characteristic energy scale* as the apparent change in kinetic energy of the NS caused by the glitch,

$$\Delta E_c \equiv 4\pi^2 I f_{\text{rot}} \Delta f_{\text{rot}}, \quad (3)$$

where I is the NS moment of inertia. Inserting values for f_{rot} and Δf_{rot} from Table 1 and a fiducial $I = 10^{45}$ g cm 2 , we obtain $\Delta E_c \approx 1.15 \times 10^{43}$ erg.

The relation of this energy to actual radiated GW energies is highly model-dependent; see Yim et al. (2024) for a summary of the literature. In particular, the energy transferred to the f-mode could be much smaller than ΔE_c if the timescale

associated with the glitch trigger (possibly crust cracking, or vortex unpinning) is much longer than the $\mathcal{O} \sim 1$ ms timescale characteristic of an f-mode oscillation, as discussed in the context of magnetar flares in Levin & van Hoven (2011). It is also possible that other modes, with oscillation periods closer to the glitch trigger timescale, might be more efficiently excited, e.g., the torsional crustal oscillations considered in the context of Vela glitches in Bransgrove et al. (2020); such oscillations are relatively weak GW emitters.

With these caveats in mind, we will nevertheless choose f-modes as the reference model for evaluating the performance of our (unmodeled) searches, and use ΔE_c as a useful energy scale to compare our results against. For a signal starting at

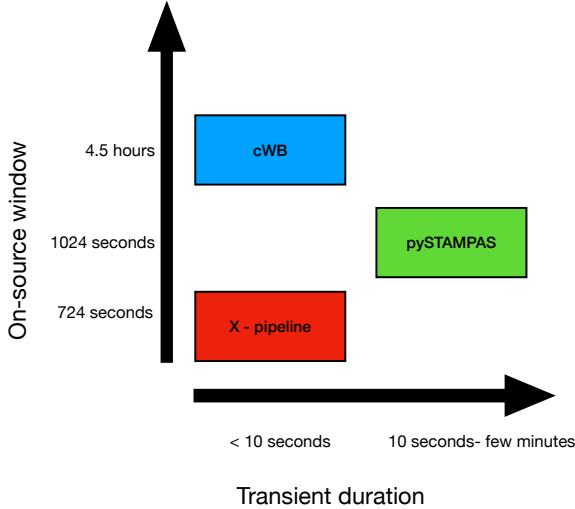


Figure 3. This schematic summarizes the burst-like transient searches performed for this paper as a function of the duration of anticipated signals and the on-source windows used by each. The on-source window time for **X-pipeline** is 1σ uncertainty around the estimated glitch arrival time using the timing model, the on-source window for **cWB** is the broadest window in radio observation within which the glitch occurred. The choice for **PySTAMPAS** is based on analysis requirements and longer duration of the signal.

t_0 , the f-mode signal can be written as (Abadie et al. 2011)

$$h_+(t) = A_+ \cos(2\pi f_{\text{gw}}(t - t_0) + \phi_0) e^{-(t-t_0)/\tau}, \quad (4a)$$

$$h_\times(t) = A_\times \sin(2\pi f_{\text{gw}}(t - t_0) + \phi_0) e^{-(t-t_0)/\tau}. \quad (4b)$$

The amplitudes A_+ and A_\times depend on the NS inclination angle ι as $1 + \cos^2(\iota)$ and $2 \cos(\iota)$. τ is the damping time, and ϕ_0 is the initial phase. The frequency and damping time depend on the equation of state of NS matter, with ranges of $\sim [1, 4]$ kHz and $\lesssim 1$ s for most realistic models (see, e.g., Andersson & Kokkotas 1998; Wilson & Ho 2024). These ranges are covered by the searches described below.

The actual emission from a glitching NS can be more complex than the simplistic assumptions made here. To mitigate this, we have employed three pipelines that do not depend explicitly on the morphology of the incoming GW signal. The damped sinusoid model of Equation (4) is then used for sensitivity measures and computing upper limits.

4.2. coherent WaveBurst (**cWB**)

cWB (Klimenko et al. 2016; Drago et al. 2021; Martini et al. 2025) is a GW search and reconstruction algorithm based on maximum-likelihood statistics that uses the excess signal power in the detector network and does not rely on a precise signal model of the incoming GWs. The time series strain data are transformed into a multi-resolution time-frequency map using the Wilson–Daubechies–Meyer (WDM) wavelet transform (Necula et al. 2012), then excess

power clusters of pixels are identified for each detector in the network, and a likelihood is computed for these selected pixels as a function of sky direction. **cWB** is used routinely for the search and reconstruction of GW transients in LVK data (Abbott et al. 2021a, 2023). The version of the pipeline used here is the same as in the O4a all-sky short-duration search (Abac et al. 2025b). **cWB** is sensitive to signals expected from NS glitches, such as f-modes (Lopez et al. 2022), and was recently used for all-sky burst searches (Abbott et al. 2021a, 2023; Abac et al. 2025a,b).

The **cWB** analysis searches for short-duration transient signals (less than 10 seconds) in the frequency range of 512 Hz to 4000 Hz. It covers a time window of 4.5 hours around the estimated glitch time (at Earth, rounded to the nearest integer, corresponding to 1398458994 GPS seconds), addressing the broadest possible uncertainty from radio observations. The analysis is restricted to sky directions within 1 square degree centered around the Vela pulsar. For the 4.5 hours of on-source time window we perform time slides to estimate the background (Was et al. 2010). The total background generated is 6.6 years, allowing false-alarm probability estimates at the level of 7.5×10^{-5} . The search does not find any triggers in the on-source time window above the thresholds applied within the pipeline.

4.3. *PySTAMPAS*

PySTAMPAS (Macquet et al. 2021) is an unmodeled coherent search pipeline based on the cross-correlation of multi-resolution time-frequency maps and targeting long duration GW transient signals. The analysis uses the same version of the pipeline as in the O4a all-sky long-duration bursts search (Abac et al. 2025a). The pipeline splits data into 512 s segments and generates signal-to-noise-ratio spectrograms for each detector using Short-Time Fourier Transforms (STFTs) over short windows (0.5, 1, 2, and 4 seconds, each Hann-windowed with 50% overlap) to form multi-resolution spectrograms covering 512–4020 Hz, with the power spectral density (PSD) in each frequency bin estimated as the median of the squared STFT magnitudes. A seed-based clustering algorithm then identifies candidate triggers, which are cross-correlated with the corresponding spectrogram from the second detector to compute a coherent detection statistic. The search is performed considering the Vela sky position given in Table 1. The on-source window is ± 512 s long centered at the glitch time 1398458994, to ensure covering the full signal power also for longer-duration signals of 50–100 seconds. We consider two days of data around the glitch time excluding the on-source window to estimate the background distribution of noise triggers by 25 time-shifts of the time-frequency maps by more than 1000 s to effectively simulate 31 days of background data. This procedure establishes a 3σ detection threshold corresponding to

a false-alarm rate (FAR) of 3.3×10^{-6} Hz. The search does not find any event in the on-source time window below this FAR, the loudest event has FAR more than an order of magnitude higher corresponding to p-value of around 0.6 which is compatible with Gaussian noise fluctuations.

4.4. X-pipeline

The X-pipeline is an un-modeled coherent search pipeline where the data from each detector is coherently combined in the multi-resolution time-frequency domain. The brightest 1% of pixels are selected along with nearby pixels that form a cluster and then assigned likelihood-based ranking statistics (Sutton et al. 2010; Was et al. 2012). Spurious noise triggers are rejected by applying vetoes to the coherent and incoherent power across the detector network.

The search performed for this paper was done on an on-source time window of ± 362 seconds centered at GPS time 1398458994, representing 1σ uncertainty in the glitch time estimate. The frequency band of the analysis is 200 Hz to 4000 Hz. The off-source time window used for background computation is ± 1.5 hours around the glitch time. The loudest event found in the analysis has p-value of 0.8 which is compatible with Gaussian noise fluctuations.

4.5. Search Sensitivity

While the burst searches are independent of a specific GW waveform, to assess their sensitivity we inject simulated damped sinusoid waveforms, as a simple stand-in model for emission from f -modes, into the detector data around the time of the glitch. These injections are fixed to the sky direction and inclination of the Vela pulsar. Since the NS equation of state and the precise glitch mechanism are not known, we also use a range of central frequencies and damping times for the damped sinusoids. We list the injection sets used for each pipeline in Table 2. These cover different timescales due to the different search ranges as previously summarized in Figure 3 and, for practical reasons, somewhat different but overlapping frequency ranges.

We then vary the amplitude of each damped sinusoid to obtain an efficiency curve for each pipeline, ranging from 0 to 100%. We fix a detection threshold of 3σ for each pipeline, which corresponds to inverse false-alarm rates of 0.2 years, 0.008 years and 0.009 years for cWB, X-pipeline and PySTAMPAS, respectively due to differences in the on-source time window. The upper limits apply only to signals during coincident data availability from the LIGO detectors, similar to choices made in previous works (e.g., the search for GWs from SN2023ixf Abac et al. 2025d).

We present the results in terms of the quantity

$$h_{\text{rss}} = \sqrt{\int_{-\infty}^{\infty} (h_+^2(t) + h_{\times}^2(t)) dt} \quad (5)$$

Table 2. Parameters of the damped sinusoid injection sets used by each pipeline

Pipeline	f_{gw} (Hz)	τ (s)
cWB	650, 1100,	0.01, 0.1, 0.5
	2100, 3900	
X-pipeline	290, 650, 1100,	0.01, 0.1, 0.5
	2100, 3900	
PySTAMPAS	800, 1050, 1300,	10
	1600, 2000, 2200,	
	2500, 2750, 3000, 3250	

where $h_+(t)$ and $h_{\times}(t)$ are the amplitudes for the two GW polarization. The Vela pulsar's inclination angle dictates the relative content of h_+ and h_{\times} of the waveform. We perform a set of injections over a range of amplitudes of $[10^{-23}, 10^{-21}]$ and recover them with each of the search pipelines, then measure the h_{rss} needed to achieve the detection efficiency 90%. We show this $h_{\text{rss}}^{90\%}$, at a 3σ detection threshold and as a function of GW frequency, for all pipelines and various damping times in Figure 4.

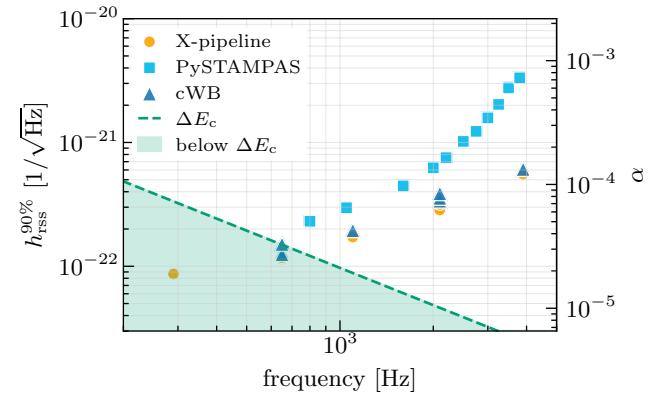


Figure 4. $h_{\text{rss}}^{90\%}$ at a detection threshold of 3σ (p-value 10^{-3}) for the three GW burst searches (circle, square and triangle markers), shown against the frequency of damped sinusoids signals. For cWB and X-pipeline, multiple markers at fixed frequencies denote different damping times (see Table 2), which do not effect the sensitivity drastically. Strains corresponding to the characteristic energy ΔE_c of the glitch from Equation (3) are shown as the dashed green line and GW emission would be consistent with it in the shaded region below. The right-hand vertical axis shows the r-mode amplitude α that would correspond to a given h_{rss} level. See Section 6.1 for interpretation of these results in terms of ΔE_c and α .

The results are different due to both methodological aspects and the different parameter spaces, as summarized in Figure 3. PySTAMPAS probes longer-duration signals, which are more difficult to detect with un-modeled methods. The difference between X-pipeline and cWB can be

in part explained by X-pipeline’s tighter on-source window. Even with the same false-alarm rate threshold between the pipelines, there would still be variations due to the antennae pattern functions and noise properties changing over the different time windows, hence a fair comparison cannot be made and is also out-of-scope of this work.

In the same figure we have also shown the h_{rss} corresponding to GWs with the full characteristic glitch energy from Equation (3); see Section 6.1 for its derivation and interpretation. The X-pipeline and cWB results compare more favorably with the available energy budget below 1 kHz than at higher frequencies, while for PySTAMPAS only injections at higher frequencies were done.

In the O4a all-sky burst search (Abac et al. 2025b), the minimum relative glitch size for Vela-like pulsars to be detected was around $\Delta f_{\text{rot}}/f_{\text{rot}} \approx 2 \times 10^{-5}$, which is an order of magnitude larger than the Δf_{gl} of this glitch. But the search sensitivity in that paper was computed with injections above 2000 Hz, which along with all-sky and all-time averaging makes those results not directly comparable to the ones in this paper.

5. SEARCHES FOR QUASI-MONOCHROMATIC LONG TRANSIENT GWs

Quasi-monochromatic long-duration GW signals at frequencies coupled to the pulsar’s rotation may be triggered by glitches (Prix et al. 2011) and can be searched for with methods derived from those for CWs; see Riles (2023) for a review. The first such searches were done on O2 data (Keitel et al. 2019), including the 2016 Vela glitch (Palfreyman et al. 2018) and a glitch of the Crab pulsar, and on O3 data (Abbott et al. 2022a; Modafferi et al. 2021), covering six glitching pulsars but with no Vela glitch during that run. Detection prospects for this type of GW signals were studied in Moragues et al. (2023), indicating that Vela glitches from O4 on would allow for physical constraints below the indirect energy estimate previously introduced by Prix et al. (2011), equivalent to our Equation (3). Yim et al. (2024) also identified Vela as a priority search target.

In this section, we first discuss emission models for this type of signal. We then present several analyses that differ in the strictness of their model assumptions: two fully-coherent and two semi-coherent search methods, each with different allowances for the degree of mismatch between the GW and EM phase evolution and varying flexibility in signal start-times and durations. The motivation is to probe as deep as possible into the noise floor for the strictest model assumptions, while also maintaining chances to detect sufficiently strong signals with different parameters or time evolution.

As there are no significant detection candidates from either search, we report upper limits on GW strain from all of them, which are compared in Section 5.6 and Figure 5. More

detailed results from each method are contained in the appendices. Detailed interpretation is left to Section 6.

5.1. Emission Models

Models for transient CW-like emission following a pulsar glitch, as summarized by Haskell & Jones (2024) and Yim et al. (2024), can be divided into two classes: GW emission caused by motions of the fluid core, or by quadrupolar deformations supported by the solid outer crust.

For fluid motion in the core, the leading model involves *Ekman pumping*, where the sudden spin-up of the crust induces a meridional flow in the interior fluid (van Eysden & Melatos 2008), which may be non-axisymmetric. Calculations by Bennett et al. (2010) and Singh (2017) found emission via both the mass and current quadrupoles and at both the f_{rot} and $2f_{\text{rot}}$ harmonics. The GW amplitude depends on the NS equation of state via the matter compressibility.

For asymmetries in the crust, the key idea is that, for any combination of elastic fracture and superfluid unpinning causing the glitch, the process happens in a non-axisymmetric way, so that a large-scale (quadrupolar) deformation is produced in the solid outer crust, which then relaxes away on some (poorly constrained) viscous timescale of the crust. This relaxation may leave an imprint in the post-glitch spin evolution, as exploited in the model of Yim & Jones (2020), which attributes the exponentially-decaying excess spin-down seen after most glitches to the formation of a transient mountain, causing GW emission that provides the increased spin-down torque.

We can again use the characteristic glitch energy scale of Equation (3) to describe such transient-CW emission, but inserting a dimensionless parameter Q describing how much of this energy is radiated in GWs:

$$\Delta E_{\text{c, tCW}} = 4\pi^2 I f_{\text{rot}} \Delta f_{\text{rot}} Q. \quad (6)$$

In the Yim & Jones (2020) transient mountain model, only the part of Δf_{rot} that decays away in the post-glitch relaxation contributes to GW emission, at $2f_{\text{rot}}$. In this case, the fraction of energy corresponds to the *healing parameter*

$$Q = \frac{\Delta f_{\text{rot}}^d}{\Delta f_{\text{rot}}} \quad (7)$$

as usually defined in pulsar timing (Antonopoulou et al. 2022). For the 2024 Vela glitch, Δf_{rot}^d can be taken as the sum of the three decay components identified in radio pulsar timing (Section 2). Using the parameters in Table 1, we find $Q \approx 0.017$, i.e., only about 1.7% of the characteristic glitch energy would be radiated in GWs in this scenario.

More energetic emission mechanisms with $Q \sim 1$ have been speculated upon (see section II.C. of Prix et al. 2011 and Table 1 of Yim et al. 2024). However, these scenarios

lack the detailed emission mechanisms of the Ekman pumping or Yim & Jones decaying mountain models, reflecting our uncertainty in the nature of transient GW emission.

Furthermore, if the NS is excited into *free precession* following the glitch, a CW-like signal would be emitted, decaying on a timescale determined by the strength of dissipative processes within the star (Zimmermann & Szedenits 1979; Jones & Andersson 2002). The signal would contain harmonics at (or close) to both f_{rot} and $2f_{\text{rot}}$, with the lower harmonic dominating in the case of a small angular misalignment between the NS symmetry axis and its (fixed) angular momentum vector. Such precessional motion would also leave an imprint on the radio pulsar signal, modulating the frequency and pulse profile (Jones & Andersson 2001). No such modulations have been reported, although they may occur with too small an amplitude and/or on too long a timescale to be apparent.

Given the wide range of long-transient GW emission mechanisms and strengths found in the literature, we carried out a number of relatively agnostic searches, as described below. All allow for a wide range of signal durations, and some also for emission at both the f_{rot} and $2f_{\text{rot}}$ harmonics. Additionally, we consider GW amplitude evolution profiles that are either constant, i.e., set by a rectangular window function

$$h(t; \lambda, \mathcal{A}, \mathcal{T}) = \begin{cases} h(t; \lambda, \mathcal{A}) & \text{if } t_0 \leq t < t_0 + \tau; \\ 0 & \text{otherwise,} \end{cases} \quad (8)$$

a single exponentially decaying window function,

$$h(t; \lambda, \mathcal{A}, \mathcal{T}) = \begin{cases} e^{-(t-t_0)/\tau} h(t; \lambda, \mathcal{A}) & \text{if } t \geq t_0; \\ 0 & \text{otherwise,} \end{cases} \quad (9)$$

or a combination of three exponentials, following Equation (2). Here, $h(t; \lambda, \mathcal{A})$ would be a standard persistent CW signal (Jaranowski et al. 1998; Riles 2023) with two polarization components, described by amplitude parameters \mathcal{A} (including a constant characteristic strain amplitude h_0 , inclination ι , polarization angle ψ , and an initial phase) and phase-evolution parameters λ (frequency, frequency derivatives, and sky position); and $\mathcal{T} = \{t_0, \tau\}$ are the transient window parameters: t_0 is the start of the signal window and τ is the rectangular window length or exponential decay constant. In the following, we also define for convenience $h_0(t) \equiv e^{-(t-t_0)/\tau} h_0$ in the exponentially decaying case, with $h_0(t = t_0) = h_0$.

All the searches and upper limits presented here are set up to correctly take into account the gaps in GW data (as seen in Figure 2), as is standard for CW search methods.

5.2. Time-domain Bayesian search (CWINPy)

The time-domain Bayesian method (Dupuis & Woan 2005; Pitkin et al. 2017) has been widely used for targeted CW

searches (e.g., Abbott et al. 2022a,b; Abac et al. 2025e), where the signal amplitude is assumed as constant over the entire observation period. The implementation in the CWINPy package (Pitkin 2022) also allows for transient time-dependent profiles as per Equations (8) or (9), though all signals considered will have fixed start time (chosen here as $t_0 = T_{\text{gl}}$) and only vary in τ (from 1 hr to 120 d).

To make the inference computationally tractable, the data are massively compressed (Dupuis & Woan 2005): First, they are heterodyned to account for the motion of the detectors with respect to the source as well as any intrinsic phase evolution, fixing to the best-fit parameters from Table 1. These heterodyned data are then lowpass-filtered and down-sampled. CWINPy then uses the bilby (Ashton et al. 2019) and dynesty (Speagle 2020) packages for nested sampling inference to produce a Bayesian evidence value and posterior probability distribution for the unknown source parameters, given the compressed data.

Eight CWINPy analyses were performed, toggling three different options: (i) a rectangular or exponential amplitude profile; (ii) unconstrained ι and ψ or observationally motivated constraints on these angles from modeling of the pulsar wind nebula (Ng & Romani 2008); (iii) emission only from the $l = m = 2$ mass quadrupole modes near $2f_{\text{rot}}$, or allowing for both the $l = 2, m = 1, 2$ modes, with emission at both the f_{rot} and $2f_{\text{rot}}$ (Jones 2010; Abbott et al. 2019) harmonics. See Appendix A for a full description of these options, other prior choices, and results.

No evidence for coherent signals across H1 and L1 was found. The posterior probability distributions on signal amplitudes have been used to define 95% credible upper limits as a function of the binned signal duration.

5.3. Transient \mathcal{F} -statistic search

This method (Prix et al. 2011) builds on the standard matched-filter CW maximum-likelihood \mathcal{F} -statistic (Jaranowski et al. 1998) in its LALSuite (LVK 2025) frequency-domain implementation. It has already been used in O2 and O3 post-glitch GW searches (Keitel et al. 2019; Abbott et al. 2022a; Modafferi et al. 2021, 2023).

We use 1800 s short Fourier transforms (SFTs) (Allen et al. 2024) produced from time-domain gated data (Davis et al. 2025), combined coherently across both LIGO detectors (Cutler & Schutz 2005) and the full observing time. The detection statistic is a Bayes factor $\mathcal{B}_{tS/G}$ marginalized over start times and durations (Prix et al. 2011), under the two amplitude profiles from Equations (8) and (9), and evaluated over a template bank in f_{rot} , \dot{f}_{rot} , \ddot{f}_{rot} , and $\ddot{\ddot{f}}_{\text{rot}}$. Unlike CWINPy, this search does not use inclination and polarization constraints. (They can be incorporated into an \mathcal{F} -statistic framework as per Jaranowski & Królak (2010) but this has not yet been implemented for transients.)

We allow for signal start times t_0 within ± 1 day of T_{gl} and duration parameters τ up to 120 days, covering data from 28 April 2024 until 28 August 2024. The search ranges are given in Table 3 of Appendix B. Rectangular-window (constant amplitude) signals can be searched for efficiently on CPUs with code in `LALSuite` (LVK 2025), while exponentially decaying signals require the GPU implementation in `PyFstat` (Keitel & Ashton 2018; Keitel et al. 2021). Both are called via the `tCWip` (“transient CW investigation pipeline”) package (Keitel et al. 2025).

We find no statistically significant candidates above a threshold set to the expected loudest outlier from an empirical background estimation (Tenorio et al. 2022a). We obtain upper limits on the initial strain amplitude h_0 via simulated signal injection and recovery. See Appendix B for details on the method, upper limits procedure and full results.

5.4. Weighted peakmap search

As a complementary approach to the two coherent searches, we also run a semi-coherent search based on weighted peakmaps (WPM). This uses time-domain band-sampled data (BSD) files (Piccinni et al. 2019), including a glitch cleaning step (Astone et al. 2005). The search starts 2000 seconds before the glitch time and covers $T_{\text{tot}} = 101.7$ days. Given the significant ASD difference of the L1 and H1 detectors, only L1 data are used. Data over a 0.2 Hz frequency band around the putative signal frequency are extracted from the BSDs and heterodyned (Abbott et al. 2022b) to remove the Doppler modulation induced by the Earth’s motion and the Vela pulsar’s secular spin-down, using values measured immediately after the glitch, see Table 4. The corrected data are divided in segments of duration T_{seg} (see Appendix C), and from each an equalized spectrum is computed. Local maxima above a given threshold are selected, creating a so-called *peakmap* (Astone et al. 2014).

We then search for peaks in the projection of the peakmap on the frequency axis, as expected for sufficiently strong signals. Peaks are weighted by their amplitude before projection, increasing the contribution of peaks from the early part of the signal, which is expected to be stronger. The detection statistic is the critical ratio $CR = (\mathcal{P} - \mu_{\mathcal{P}})/\sigma_{\mathcal{P}}$, where \mathcal{P} is the sum of the weighted peaks and $\mu_{\mathcal{P}}, \sigma_{\mathcal{P}}$ are the corresponding projection mean value and standard deviation.

This procedure is independent of the assumed signal model. But by assuming a specific signal model we can make an “optimal” choice of the observation window $T_{\text{obs}} \leq T_{\text{tot}}$ and segment duration T_{seg} – see Appendix C and the resulting choices in Table 5 for the seven considered search setups.

For each of these setups, we select the ten most significant outliers as those with the highest CR from each of ten 0.02 Hz sub-bands. Those with false-alarm probability (taking into account the trials factor) larger than 0.01 are dis-

carded. Three outliers remain: one corresponding to $\tau = 8$ days with frequency 22.455 Hz ($CR = 6.65$), one to $\tau = 30$ days with frequency 22.5011 Hz ($CR = 7.55$) and one to $\tau = 1000$ days with the same frequency ($CR = 5.15$). They are subjected to a follow-up procedure based on different methods, see Appendix C.2. While the last two are associated to an instrumental spectral line, the first one is more difficult to interpret, see Appendix C.2. As it is statistically marginal, we do not consider it a promising detection candidate, and also proceed for this method to compute upper limits.

The upper limits are placed in terms of the initial signal strain amplitude via simulated signal injections. For the polarization parameters ψ and ι we use both the restricted prior described for `CWInPy` in Appendix A and an unrestricted uniform distribution. See Appendix C for details, with Figure 12 and Table 5 including the full results.

5.5. Hidden Markov model search

The fourth search for transient quasi-monochromatic GWs combines the semi-coherent \mathcal{F} -statistic (Jaranowski et al. 1998; LVK 2025) with hidden Markov model (HMM) tracking (Suvorova et al. 2016; Sun et al. 2018) at both once and twice the Vela spin frequency (Sun et al. 2019). It uses the same 1800 s SFTs as the transient \mathcal{F} -statistic method.

Instead of a fully coherent analysis, the HMM pipeline divides the data into segments of duration T_{coh} and evaluates the standard \mathcal{F} -statistic over a template grid coherently in each segment (without explicitly applying time-windowed search templates), which are then recombined incoherently.

The HMM signal model allows for a drift in frequency between each segment, i.e., for stochastic spin-wandering or other deviations from the standard Taylor phase model of Equation (1). We employ the Viterbi algorithm (Viterbi 1967) to determine the optimal frequency path for both the single harmonic ($2f_{\text{rot}}$) and dual harmonic (f_{rot} and $2f_{\text{rot}}$) scenarios. The detection statistic is the log-likelihood of the optimal path normalized by the number of segments, $\bar{\mathcal{L}} \equiv \mathcal{L}/N_T$.

The configuration of the HMM search is described in Table 7 in Appendix D. It covers a 0.05 Hz band centered on $f_{\text{gw}} = f_{\text{rot}}$ and 0.1 Hz band centered on $f_{\text{gw}} = 2f_{\text{rot}}$, with frequency resolutions of $1/(4T_{\text{coh}})$ and $1/(2T_{\text{coh}})$, respectively. The observation window covers $T_{\text{obs}} = 90$ days. For added flexibility, we repeat the HMM analysis with three values of T_{coh} , each corresponding to one of the three glitch recovery timescales. Shorter T_{coh} reduces sensitivity, but can track a more rapidly fluctuating signal.

Empirical detection thresholds are set using two methods (see Appendix D), based on either repeating the search on numerous realizations of Gaussian noise data, or repeating the search at randomly chosen off-target sky positions.

The full results of the HMM search are compiled in Table 8, where we list the max- $\bar{\mathcal{L}}$ statistics obtained with single or dual harmonic tracking for each value of T_{coh} . Our analysis does not recover any outliers with false-alarm probability below 1%. We therefore find no statistically significant outliers for the HMM search.

We set upper limits by injecting simulated transient signals with varying durations τ into the detector data, and then recovering the injection using the same HMM search configurations, with single or dual harmonic tracking. Upper limit results for the single harmonic emission ($2f_{\text{rot}}$ only) scenario with $T_{\text{coh}} = 9$ hr are included in Figure 5, with the full set of upper limits shown in Appendix D and Figure 14.

5.6. Comparison of upper limits results

In Figure 5 we compare a selection of upper limits from each of the four methods, focusing on signals at twice the pulsar's rotation frequency and with exponential amplitude decay. The results were chosen to span the range of assumptions of varying strictness covered by the set of pipelines: the strictest upper limits for such signals are obtained by the fully-coherent targeted `CWInPy` search with constrained orientation angles, while the semi-coherent HMM search with the least prior assumptions provides the least strict constraints. The two other methods and other configuration choices provide intermediate strictness of both assumptions and resulting constraints. The figure also contains comparisons with theoretical expectations on possible post-glitch GW strains, which will be discussed in the following section.

These limits concern arbitrary exponentially decaying CW-like transients within the search ranges, and at least for the \mathcal{F} -statistic method it is known that the exact amplitude evolution only has a small effect on its sensitivity (Prix et al. 2011; Keitel et al. 2019). In addition, from the WPM method we also report a dedicated upper limit for GW amplitudes following the three-component frequency decay observed in the radio, for which we find $h_0^{95\%} = 4.5 \times 10^{-24}$ assuming uniform priors on the polarization parameters and 4.7×10^{-24} for restricted priors. (See also Table 5 in the appendix.)

How deep a search has probed below the detector noise floor can be stated as a sensitivity depth (Behnke et al. 2015; Dreissigacker et al. 2018; Wette 2023) $\mathcal{D} = \sqrt{S_n(f)/h_0}$, where $\sqrt{S_n}$ is the ASD suitably averaged over the observation time. These searches have probed below the average four-month H1+L1 ASD by \mathcal{D} factors from $4 \text{ Hz}^{-1/2}$ for the most flexible HMM search (at $\tau = 1$ d) to $170 \text{ Hz}^{-1/2}$ for the most restrictive `CWInPy` analysis (at $\tau = 120$ d).

6. INTERPRETING THE OBSERVATIONAL CONSTRAINTS

We now provide some interpretation of our results. In Section 6.1 we discuss what conclusions can be drawn given the

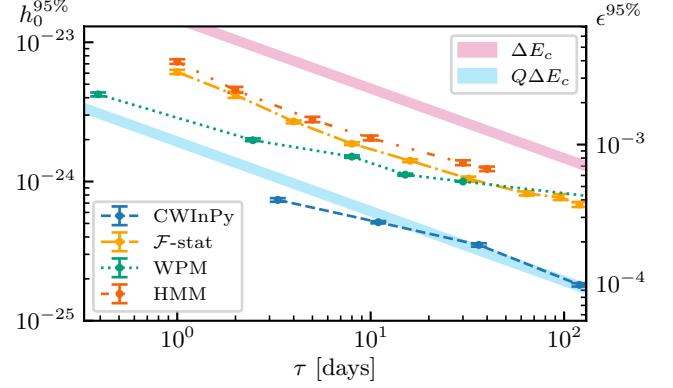


Figure 5. Selected upper limits at 95% confidence from the four CW-like post-glitch searches, in terms of initial strain amplitude (left-hand axis) and NS ellipticity ϵ (right-hand axis), both defined at the start of signals with exponentially decay parameters τ . The results were chosen to illustrate how upper limits scale with the strictness of prior assumptions built into each search configuration. The `CWInPy` results have the strictest constraints (fixed frequency evolution and constrained orientation angles), while the three other results are for narrow-band searches, with \mathcal{F} -statistic results marginalized over all durations, WPM results for search configurations optimized at each τ point, and HMM results for $T_{\text{coh}} = 9$ h.

For comparison, GW emission at the characteristic energy scale ΔE_c from Equation (3) is shown as the magenta shaded band near the top, while the sky blue shaded band further down shows the emission if the GW energy output is suppressed by a factor $Q \approx 0.017$, as in the Yim & Jones (2020) transient mountain model.

lack of detection of a burst-like signal. In Section 6.2 we discuss the same for the non-detection of long-duration transient quasi-monochromatic signals. In Section 6.3 we discuss the combined constraints from both on GW emission models and NS properties.

Before looking at the separate burst and long-transient cases, it is useful to assemble some formulae relevant to both. Despite searching for signals over very different frequencies and durations, both types of search cover signals of fixed source frequency and of finite duration. Hence, we can also describe the burst analyses from Section 4 (for the case of f-mode or similar signals) with the same notation as for CW-like signals from Section 5, using the explicit $h_0(t)$ and converting to the average h_{rss} where needed. The total GW energy in such a signal is then given by (see e.g. equation (3) of Prix et al. 2011):

$$E_{\text{gw}} = \frac{2\pi^2}{5} \frac{c^3}{G} f_{\text{gw}}^2 d^2 \int_{t_0}^{t_0 + T_{\text{obs}}} h_0^2(t) dt, \quad (10)$$

where T_{obs} is the observation time. This applies regardless of the time variation of the signal amplitude $h_0(t)$. The corresponding optimal signal-to-noise ratio is

$$\rho_0^2 = \int_{t_0}^{t_0 + T_{\text{obs}}} \frac{B(t; \hat{n}, \iota, \psi) h_0^2(t)}{S_h(t, f_{\text{gw}})} dt, \quad (11)$$

where $S_h(t, f_{\text{gw}})$ is the power spectral density of the detector noise and $B(t; \hat{n}, \iota, \psi)$ summarizes location- and orientation-dependent factors (for details see Jaranowski et al. 1998; Prix et al. 2011; Prix 2015).

For a simple exponential decay on timescale τ , the amplitudes h_{rss} and h_0 are related by

$$h_{\text{rss}} = h_0 \sqrt{\frac{\tau}{2}}, \quad (12)$$

as can easily be verified by assuming an exponential time dependence for the integrand in Equation (5).

We will now use these formulae to interpret the results from Sections 4 and 5, and compare with the characteristic energy scale of Equation (3) and additional model-specific constraints.

6.1. Interpretation: burst results

The results of our burst searches were presented in Figure 4. For the decaying exponentials assumed in the injections to set these limits, Equation (11) reduces to $\rho_0^2 \propto h_{\text{rss}}^2 / S_h(f_{\text{gw}})$, defining $S_h(f_{\text{gw}})$ as a time average over $S_h(t, f_{\text{gw}})$. Hence, at fixed signal-to-noise ratio, the minimum detectable h_{rss} scales as $\propto \sqrt{S_h(f_{\text{gw}})}$. This accounts for the limits weakening with increasing frequency, reflecting the slope of the detector noise curve as seen in Figure 2.

To interpret the astrophysical significance of the results, we can similarly re-write Equation (10) to give

$$E_{\text{gw}} = \frac{2\pi^2}{5} \frac{c^3}{G} f_{\text{gw}}^2 d^2 h_{\text{rss}}^2. \quad (13)$$

Setting E_{gw} equal to the characteristic energy of Equation (3) then gives the corresponding h_{rss} as a function of f_{gw} , with scaling $h_{\text{rss}} \propto 1/f_{\text{gw}}$. Inserting values for the 29 April 2024 Vela glitch, we obtain $h_{\text{rss}}(f_{\text{gw}}) \approx 9.7 \times 10^{-20} \text{ Hz}^{-1/2} / (f_{\text{gw}}/\text{Hz})$, which is plotted as the dashed green line in Figure 4.

The search upper limits constrain the actual GW energy to be less than this at sufficiently low frequencies. For instance, for the X-pipeline this characteristic energy scale is surpassed for frequencies below about 800 Hz, and at the lowest injection frequency of 290 Hz, h_{rss} is constrained to a factor of ≈ 3 below the characteristic scale. Given the scaling $E_{\text{gw}} \sim h_{\text{rss}}^2$ of Equation (13), this corresponds to a factor of ≈ 9 in radiated energy. However, emission at such low frequencies would not correspond to the f-modes considered most plausible to be excited by pulsar glitches. For example, see Wilson & Ho (2024) where f-mode emission was studied for a catalog of 15 equations of state (eight for standard NSs and seven for quark stars).

Our upper limits on h_{rss} can be converted into upper limits on a (dimensionless) amplitude α of the f-mode excitation, such that a point on the surface of the NS oscillates

about its equilibrium position by a distance $\sim \alpha R$, with R the stellar radius. This conversion depends on the NS mass M and equation of state. To provide a quantitative example of this, we will make use of the analytic results given in Yim & Jones (2023) for uniform density stars. Combining equations (5) and (10) of Yim & Jones (2023), and making use of our Equation (12), one obtains

$$\alpha = \left(\frac{2\pi}{15} \right)^{1/2} \left(\frac{c^3}{G} \right)^{1/2} \frac{h_{\text{rss}} d}{M^{1/2} R}. \quad (14)$$

Inserting $M = 1.4 M_{\odot}$ and $R = 10^6$ cm, and the distance for Vela, we find

$$\alpha \approx 7 \times 10^{-5} \left(\frac{h_{\text{rss}}}{3 \times 10^{-22} \text{ Hz}^{-1/2}} \right), \quad (15)$$

where we have chosen the scaling in h_{rss} based on the upper limits on h_{rss} of Figure 4, at a frequency of about 1.5 kHz, typical for an f-mode. Roughly speaking, we have constrained the oscillation amplitude αR of a particle on the Vela's surface to be no larger than about 70 cm for a typical f-mode frequency. However, at these frequencies our upper limits on h_{rss} fall short of the values corresponding to the characteristic energy by about a factor of ~ 4 .

6.2. Interpretation: long transient results

As discussed in Section 5.1, there are two main mechanisms for producing CW-like transients following a glitch. The first we discussed was Ekman pumping. Using equation (55) of Bennett et al. (2010) and the parameters for the 29 April 2024 Vela glitch, we obtain an estimate of the signal strength of $h_0 \approx 7 \times 10^{-26}$. While this estimate is conservative, mainly due to uncertainties on the lag between the components in the NS interior (which may be larger than the observed glitch size, Haskell & Melatos 2015), and unknown nuclear physics parameters that enter the model, the predicted strain is significantly lower than the upper limits on h_0 from our analyses (Figure 5), which are hence not yet sensitive enough to set constraints on this mechanism.

The second emission mechanism we discussed was the formation of a transient mountain that then relaxes away (Yim & Jones 2020). To help interpret results from Figure 5 under this mechanism, we can insert the assumption of exponential decay on a timescale τ , Equation (9), as would be expected for dissipative mountain decay, into Equation (10) for E_{gw} . If we assume $\tau \ll T_{\text{obs}}$, we obtain

$$E_{\text{gw}} = \frac{\pi^2}{5} \frac{c^3}{G} f_{\text{gw}}^2 d^2 h_0^2(0) \tau. \quad (16)$$

For a fixed E_{gw} , this shows that $h_0 \sim \tau^{-1/2}$, explaining the slope of the upper limits in the figure.

To understand the astrophysical significance of these results, we can use Equation (6) where only a fraction Q of

the characteristic glitch energy of Equation (3) goes into GW emission. This yields for the initial GW strain amplitude

$$h_0^{Q\Delta E_c} \approx \frac{1}{d} \sqrt{\frac{5G}{2c^3} \frac{I}{\tau} Q \frac{\Delta f_{\text{rot}}}{f}}. \quad (17)$$

For $Q = 1$, this corresponds to what was discussed as “indirect energy limits” in [Prix et al. \(2011\)](#); [Keitel et al. \(2019\)](#); [Abbott et al. \(2022a\)](#), and is shown (as a function of τ) as the magenta band in Figure 5. The width of the band includes uncertainties on the Vela pulsar distance and its timing parameters, but not on the NS moment of inertia which could vary by up to another factor of 3 for conventional models or 5 for extreme matter compositions ([Lattimer 2021](#); [Johnson-McDaniel & Owen 2013](#)). The observational GW upper limits from all analyses are stricter than this by factors of ~ 3 to ~ 10 . Notwithstanding the phenomenological nature of this indirect benchmark, this is a significant milestone for post-glitch transient GW analyses.

To interpret the results in terms of a more specific model, we consider the [Yim & Jones \(2020\)](#) prediction of the energy fraction Q corresponding to the healing parameter defined in Equation (7). The 29 April 2024 Vela glitch had three observed recovery components (see Equation 2), but to be consistent with the other quantities in Figure 5, we show as the sky-blue band a simplified result, approximating the logic of the model: the result of Equation (17) with $Q \approx 0.017$ and as a function of a single τ .

Most upper limits are still above this prediction, while the result from the most constrained `CWInPy` search just reaches this level of emission. So under optimal assumptions, a detection would have been just about possible under this model, depending on the assumed GW signal duration τ and the actual moment of inertia I of the Vela NS. The quantitative implications of non-detection under this model are considered in Section 6.3.

We can also recast our results in terms of the additional ellipticity ϵ imparted to the Vela pulsar right after the glitch, which is related to the GW amplitude by the equation (see e.g. [Riles 2023](#))

$$h_0 = \frac{4\pi^2 G \epsilon I f_{\text{gw}}^2}{c^4 d}. \quad (18)$$

When applied to Vela this gives

$$\epsilon \approx 5.4 \times 10^{-4} \left(\frac{h_0}{10^{-24}} \right), \quad (19)$$

interpreted as the initial transient ellipticity increase that then decays away again.

From Figure 5, we see that our limits on h_0 span the range $[2 \times 10^{-25}, 8 \times 10^{-24}]$, depending on the search and on τ , corresponding to an ellipticity range of $[10^{-4}, 3 \times 10^{-3}]$. In

an astrophysical context, these are very large values—about three orders of magnitude larger than the largest mountains that a NS crust is estimated to be able to support (see e.g. [Johnson-McDaniel & Owen 2013](#); [Glampedakis & Gualtieri 2018](#)). However, such estimates assume a *persistent* (i.e. infinitely long-lived) mountain. A realistic crust may respond *plastically* to the stresses that support such a mountain ([Baiko & Chugunov 2018](#)), so that the exponential post-glitch relaxation may represent the plastic decay of a larger asymmetry. For this reason, the upper limits we have obtained are of astrophysical interest, despite being larger than the theoretical ones for permanent mountains. For further comparison, the `CWInPy` upper limit on a persistent ellipticity of the Vela pulsar from 8 months of O4a data was $\epsilon \leq 4.7 \times 10^{-5}$, also at 95% confidence ([Abac et al. 2025e](#)).

As discussed in Section 5.1, if the glitch caused the Vela pulsar to go into transient free precession, there would be additional emission at $f_{\text{gw}} \approx f_{\text{rot}}$. Due to the steep detector noise curve at low frequencies, strain upper limits are less constraining by an order of magnitude in this regime. Quantitative details are given in appendices A and D.

6.3. Interpretation: joint analysis and NS properties

We now show how our non-detections can be used to make statements about the NS parameter ranges under which specific GW emission models can be valid. Since this is the first time that both short- and long-duration GW searches have been performed for the same pulsar glitch, we aim to illustrate the benefit of considering results from the two types of searches jointly. So we assume that both a burst signal and a longer-duration quasi-monochromatic transient were produced at the glitch. The fact that these signals were *not* detected either (i) rules out emission models that predict GW energy fractions in one of the signal channels that our searches would have detected, or (ii) places non-trivial constraints on the NS mass and radius in the sense that Vela could still have emitted GW energy as predicted by the models, but at parameters leading to reduced detectability.

Following the method presented by [Ball & Frey \(2025\)](#), we assume specific models for both short- and long-duration waveforms. For bursts, we assume a rapidly damping sinusoid, consistent with an f-mode, in the form used by [Ho et al. \(2020\)](#) and parameterized by h_{rss} , as discussed in Section 4.1. Meanwhile, the long-transient model corresponds to the decay of a transient mountain ([Yim & Jones 2020](#)), as discussed in Section 5.1, with an initial strength h_0 .

We parameterize E_{gw} in each channel as some fraction F of the energy ΔE_c from Equation (3). This gives $E_{\text{gw}} = E_{\text{CW}} + E_{\text{fmode}} = F_{\text{CW}} \Delta E_c + F_{\text{fmode}} \Delta E_c$. We employ an equation-of-state independent universal relation from [Yagi & Yunes \(2017\)](#) to connect the pulsar’s moment of inertia I to the non-rotating mass and radius. We use another

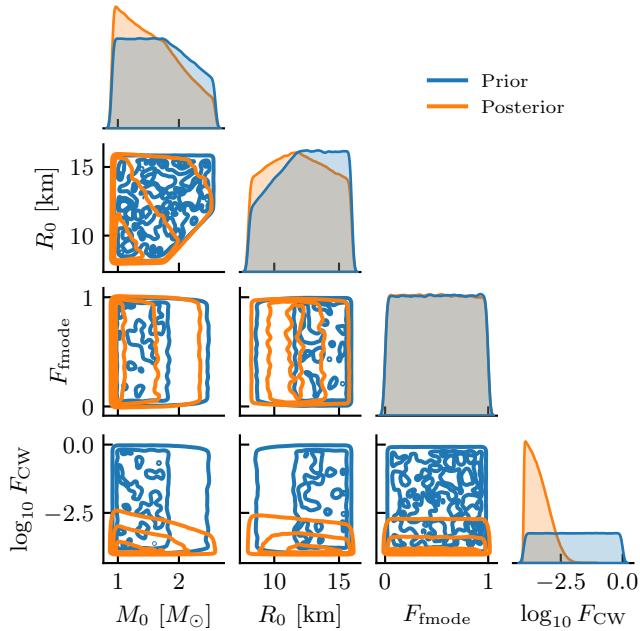


Figure 6. Prior and posterior distributions of NS mass, radius and energy fractions F_{fmode} and $\log_{10}(F_{\text{CW}})$ under the assumption that GWs from both a short-duration f-mode and a long-duration transient mountain scenario were emitted but not detected. After including the search results, there is little effect on the allowed f-mode energy in the physically motivated region $F_{\text{fmode}} < 1$ but a measurable effect on the allowed CW energy (F_{CW}). Contours show 10%, 50%, and 90% credible intervals.

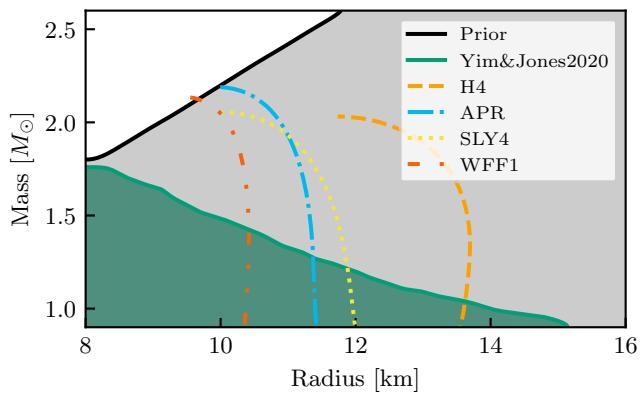


Figure 7. NS mass and radius regions for which the Vela pulsar could have emitted GWs according to the Yim & Jones (2020) model after its 2024 glitch, but without our searches detecting them (within the green shaded region at the lower left, at 90% confidence). This is compared to the assumed prior ranges (larger black contour) and to mass–radius curves for some example equations of state (Wiringa et al. 1988; Akmal et al. 1998; Douchin & Haensel 2001; Lackey et al. 2006), for which allowed NS configurations would lie along each curve. Here we fixed $F_{\text{CW}} = 0.017$ to match the healing parameter Q defined in Equation (7) and arbitrarily chose $F_{\text{fmode}} = 0.8$ (which has little effect). We see that the Yim & Jones model is only compatible with our non-detection for smaller mass and radius ranges.

set of equation-of-state independent universal relations from Pradhan et al. (2022) to connect the f-mode frequency and damping time to the non-rotating mass and radius.

We perform Monte Carlo simulations of the parameter space, drawing samples from a set of prior distributions. For each prior sample, the initial CW-like signal amplitude h_0 and the f-mode integrated strain h_{rss} are calculated. These values are compared to detection efficiency curves of a search pipeline for simulated waveforms, i.e., the probability of the search recovering a signal with some morphology at different amplitudes. The priors are described in Appendix E, along with how the likelihood function is constructed from the X-pipeline analysis and the most constrained CWInPy configuration as representatives of the two search types.

We first let F_{fmode} and F_{CW} independently vary up to a maximum of 1. The posterior distribution for mass, radius, F_{fmode} , and F_{CW} is shown in Figure 6. Consistent with the discussion above of the relative strictness of the upper limits from both types of searches in comparison to the ΔE_c energy scale, burst search limits have little to no effect on the posterior distribution, while CW search limits strongly constrain F_{CW} . From this posterior, we constrain F_{CW} to below 5.65×10^{-3} with 95% confidence.

Given the minimal effect of F_{fmode} , we then fix it to an arbitrary value of 0.8 and explore the allowed mass and radius values if we assume the Yim & Jones (2020) model correctly predicts F_{CW} , yet no detection was made. Figure 7 shows how setting F_{CW} equal to the healing parameter $Q \approx 0.017$ defined in Equation (7) requires a smaller mass and radius for the Vela pulsar if GW emission as, per this model, would have been present after this glitch. This effect is driven by the moment of inertia relation from Pradhan et al. (2022).

In addition to being dependent on the specific emission model considered and the tightest prior assumptions among the set considered in our analyses, these results are still less informative compared to other NS mass–radius limits from the binary NS merger GW170817 (Abbott et al. 2017, 2018), observations of NSs in binary systems (e.g., Taylor & Weisberg 1982; Barr et al. 2024), or NICER pulse profile modeling (e.g., Riley et al. 2021; Raaijmakers et al. 2021).

Applying this method to future glitches with improved GW detector sensitivity, a continued lack of detection could place stronger constraints on models such as Yim & Jones (2020), allow us to probe different emission scenarios, and to gain stronger benefits from a joint burst–CW analysis.

7. CONCLUSIONS

Pulsar glitches are unique probes of the structure and dynamics of NSs. The improved sensitivity of the LIGO detectors during the O4 run and the detailed timing information from the IAR and MPRO radio telescopes have now allowed us, for the first time, to place physically meaningful upper

limits on the GW emission from a glitching pulsar, constraining it to below the indirectly inferred energy scale for at least part of the relevant parameter space.

Informed by an ephemeris fit to the combined IAR–MPRO data, we have analyzed LIGO data around and after the 29 April 2024 glitch of the Vela pulsar, one of the closest and most prolific glitching pulsars known. None of the GW searches found promising detection candidates.

Our search with three unmodelled methods for GW bursts covering different signal durations from milliseconds to minutes has set constraints on the GW strain that up to 700 Hz are, for the first time, stricter than the values corresponding to the full glitch excess energy. The tightest constraint is $h_{\text{rss}} \approx 8.6 \times 10^{-23}$ at 290 Hz, a factor of ~ 3 below the characteristic energy scale.

We also searched for quasi-monochromatic long-duration transients of up to four months near twice the Vela pulsar’s rotation frequency. With four methods of varying strictness in their assumptions, we have excluded GW strain as low as a factor of 10 below the characteristic scale if all energy liberated at the glitch were converted into exponentially decaying quasi-monochromatic GWs and under the strictest assumptions, and still put constraints below that energy scale even under relaxed assumptions and phase evolution models. We also searched for similar signals from free precession at once the rotation frequency, but results are less constraining due to the steep rise of the LIGO noise curve towards lower frequencies.

Previous searches for GW bursts from the Vela pulsar (Abadie et al. 2011), all-sky GW burst searches (Abac et al. 2025a,b), and quasi-monochromatic GW transient searches for Vela and other glitching pulsars (Keitel et al. 2019; Abbott et al. 2022a; Modafferi et al. 2021, 2023) all still yielded results above the characteristic glitch energy scale. In contrast, our new results for the first time allow meaningful comparisons with specific emission models.

Under the burst scenario, we can constrain emission only at lower frequencies but not yet at a kHz or above, which most NS equations of state predict (see, e.g., Andersson & Kokkotas 1998; Wilson & Ho 2024), but where detector sensitivity is worse, primarily being limited by laser shot noise (Capote et al. 2025).

Under the “transient mountain” scenario of Yim & Jones (2020), the scale of long-duration emission is set by the transient part of the glitch’s spin-down change, which in the case of this Vela glitch is unfortunately a small fraction, and hence we only reach the sensitivity to probe this model’s predictions under the most constraining prior assumptions in one of our searches.

Combining the most constraining burst and long-duration search results, respectively, we have also illustrated how a joint analysis can constrain the NS mass and radius

ranges consistent with specific emission models under a non-detection of post-glitch GWs. However, further improvement particularly in the f-mode upper limits will be required for the full benefits of this method.

With the sensitivity of the LVK detector network further improving towards its fifth observing run (Abbott et al. 2020) and the Vela pulsar producing large glitches every two to three years, we can expect further improvements to this type of analysis in the coming years. Other glitching pulsars can also be promising targets, though the Vela pulsar stands out due to its low distance and strong glitches (Lopez et al. 2022; Moragues et al. 2023; Yim et al. 2024). Going beyond the non-detection constraints we introduced here, a positive detection of GWs from a glitching pulsar would be another important milestone in NS astrophysics.

Data products associated with this paper are released at <https://doi.org/10.5281/zenodo.17735648>.

8. ACKNOWLEDGMENTS

This material is based upon work supported by NSF’s LIGO Laboratory, which is a major facility fully funded by the National Science Foundation. The authors also gratefully acknowledge the support of the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO 600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. The authors gratefully acknowledge the Italian Istituto Nazionale di Fisica Nucleare (INFN), the French Centre National de la Recherche Scientifique (CNRS) and the Netherlands Organization for Scientific Research (NWO) for the construction and operation of the Virgo detector and the creation and support of the EGO consortium. The authors also gratefully acknowledge research support from these agencies as well as by the Council of Scientific and Industrial Research of India, the Department of Science and Technology, India, the Science & Engineering Research Board (SERB), India, the Ministry of Human Resource Development, India, the Spanish Agencia Estatal de Investigación (AEI), the Spanish Ministerio de Ciencia, Innovación y Universidades, the European Union NextGenerationEU/PRTR (PRTR-C17.I1), the ICSC - CentroNazionale di Ricerca in High Performance Computing, Big Data and Quantum Computing, funded by the European Union NextGenerationEU, the Comunitat Autònoma de les Illes Balears through the Conselleria d’Educació i Universitats, the Conselleria d’Innovació, Universitats, Ciència i Societat Digital de la Generalitat Valenciana and the CERCA Programme Generalitat de Catalunya, Spain, the Polish National Agency for Academic Exchange, the National Science Centre of Poland and the European Union - European Re-

gional Development Fund; the Foundation for Polish Science (FNP), the Polish Ministry of Science and Higher Education, the Swiss National Science Foundation (SNSF), the Russian Science Foundation, the European Commission, the European Social Funds (ESF), the European Regional Development Funds (ERDF), the Royal Society, the Scottish Funding Council, the Scottish Universities Physics Alliance, the Hungarian Scientific Research Fund (OTKA), the French Lyon Institute of Origins (LIO), the Belgian Fonds de la Recherche Scientifique (FRS-FNRS), Actions de Recherche Concertées (ARC) and Fonds Wetenschappelijk Onderzoek - Vlaanderen (FWO), Belgium, the Paris Île-de-France Region, the National Research, Development and Innovation Office of Hungary (NKFIH), the National Research Foundation of Korea, the Natural Sciences and Engineering Research Council of Canada (NSERC), the Canadian Foundation for Innovation (CFI), the Brazilian Ministry of Science, Technology, and Innovations, the International Center for Theoretical Physics South American Institute for Fundamental Research (ICTP-SAIFR), the Research Grants Council of Hong Kong, the National Natural Science Foundation of China (NSFC), the Israel Science Foundation (ISF), the US-Israel Binational Science Fund (BSF), the Leverhulme Trust, the Research Corporation, the National Science and Technology Council (NSTC), Taiwan, the United States Department of Energy, and the Kavli Foundation. The authors gratefully acknowledge the support of the NSF, STFC, INFN and CNRS for provision of computational resources. This work was supported by MEXT, the JSPS Leading-edge Research Infrastructure Program, JSPS Grant-in-Aid for Specially Promoted Research 26000005, JSPS Grant-in-Aid for Scientific Research on Innovative Areas 2402: 24103006, 24103005, and 2905: JP17H06358, JP17H06361 and JP17H06364, JSPS Core-to-Core Program A. Advanced Research Networks, JSPS Grants-in-Aid for Scientific Research (S) 17H06133 and 20H05639, JSPS Grant-in-Aid for Transformative Research Areas (A) 20A203: JP20H05854, the joint research program of the Institute for Cosmic Ray Research, University of Tokyo, the National Research Foun-

dation (NRF), the Computing Infrastructure Project of the Global Science experimental Data hub Center (GSDC) at KISTI, the Korea Astronomy and Space Science Institute (KASI), the Ministry of Science and ICT (MSIT) in Korea, Academia Sinica (AS), the AS Grid Center (ASGC) and the National Science and Technology Council (NSTC) in Taiwan under grants including the Science Vanguard Research Program, the Advanced Technology Center (ATC) of NAOJ, and the Mechanical Engineering Center of KEK. The activities at the Instituto Argentino de Radioastronomía (IAR) are supported by the national agency CONICET, the Province of Buenos Aires agency (CIC), and the National University of La Plata (UNLP).

Additional acknowledgements for support of individual authors may be found in the following document:

<https://dcc.ligo.org/LIGO-M2300033/public>. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) license to any Author Accepted Manuscript version arising. We request that citations to this article use 'A. G. Abac *et al.* (LIGO-Virgo-KAGRA Collaboration), ...' or similar phrasing, depending on journal convention. This article has been assigned document number [LIGO-P2500086](#).

Facilities: LIGO, MtPO:26m, IAR

Software: astropy (Robitaille *et al.* 2013; Price-Whelan *et al.* 2018), cWB (Klimenko *et al.* 2016; Drago *et al.* 2021; Martini *et al.* 2025), CWInPy (Pitkin 2022), distromax (Tenorio *et al.* 2022a,b), DSPSR (van Straten & Bailes 2011), gwpy (Macleod *et al.* 2021), LALSuite (LVK 2025), PyFstat (Keitel *et al.* 2021; Ashton *et al.* 2025), matplotlib (Hunter 2007), numpy (Harris *et al.* 2020), PRESTO (Ransom 2011), PSRCHIVE (Hotan *et al.* 2004), PySTAMPAS (Macquet *et al.* 2021), ROOT (Brun & Rademakers 1997), scipy (Virtanen *et al.* 2020), scikit-image (van der Walt *et al.* 2014), SWIGLAL (Wette 2020), tCWip (Keitel *et al.* 2025), TEMPO2 (Hobbs *et al.* 2006; Edwards *et al.* 2006), X-pipeline (Sutton *et al.* 2010; Was *et al.* 2012).

APPENDIX

This appendix includes additional details on the setups and results of the four CW-like long-duration post-glitch searches covered in Section 5 of the main text, as well as of the joint analysis of burst and CW results covered in Section 6.3.

A. DETAILS OF THE CWInPy ANALYSIS

As mentioned in Section 5.2, eight search configurations were used for fully-coherent targeted time-domain Bayesian analyses with the CWInPy package, corresponding to all combinations of toggling three options. These options are described here in more detail, followed by the full results.

The first option allows the GW strain amplitude to have either a rectangular or exponential decay profile. The rectangular analysis uses a uniform prior on its duration τ ranging from $\tau_{\min} = 1$ h to $\tau_{\max} = 120$ d, i.e., $p(\tau) = (\tau_{\max} - \tau_{\min})^{-1}$. The prior on the exponential profile's decay factor τ has the same range. In both cases, t_0 is fixed at the glitch epoch.

The second option either allows the orientation angle ι and polarization angle ψ of the source to be unconstrained over their physically allowable ranges or to use observationally motivated constraints on their prior ranges derived from modeling of the pulsar wind nebula (Ng & Romani 2008). The angle ι corresponds to the inclination of the rotation axis with respect to the line of sight, and ψ measures the orientation of the rotation axis in the plane perpendicular to the line of sight. For the unconstrained case, the prior distributions for both are uniform over the surface of a sphere, such that $p(\iota) = \sin \iota$ over the range $[0, \pi]$ and $p(\psi) = 2\pi^{-1}$ over the range $[0, \pi/2]$. When constrained, the prior distribution for ι is a double-peaked Gaussian distribution with peaks at 1.1048 and 2.0368 radians and $\sigma = 10.5$ mrad, and the prior distribution for ψ is a Gaussian distribution centered at 2.2799 radians and $\sigma = 1.5$ mrad (see Appendix B of Abbott et al. 2017).

The third option covers the cases of a search assuming emission only from the $l = m = 2$ mass quadrupole mode at only twice the rotation frequency, or a search with potential emission from both the $l = 2, m = 1, 2$ modes, with dual-harmonic emission at both once and twice the rotation frequency (Jones 2010; Abbott et al. 2019). In the latter case there is a characteristic amplitude for each mode, C_{21} and C_{22} , both of which will have the same decay profile applied to them. This also requires the input data to be processed via heterodyning at both frequencies.

In all cases, we use a broad unphysical uniform prior on the amplitude h_0 , or on both C_{21} and C_{22} for the dual-harmonic mode (see equations (1) and (2) of Abbott et al. 2022b), bounded in the range $[0, 10^{-21}]$. The upper range is chosen to be large enough that the likelihood will be negligible at this point, while otherwise allowing the likelihood to dominate the posterior. This allows to provide upper limits that are based on the data only.

Furthermore, in each search configuration, data from the H1 and L1 detectors were analyzed both coherently and incoherently (i.e., with the H1 and L1 datasets combined or analyzed separately), providing Bayesian evidence values for a coherent signal between detectors and independent signals in each detector. These evidences, along with that from assuming purely Gaussian noise in both detectors, have been combined to give an odds ratio comparing a coherent multi-detector signal to incoherent signals in each detector or noise (see, e.g. Abbott et al. 2017).

None of the results indicate more evidence for a coherent signal than for either an incoherent signal or none at all, with the highest \log_{10} Bayesian odds only reaching -3.45 from the configuration with exponential decay, unconstrained $(\iota, \psi), l = m = 2$ mode only. Full results from all configurations are included in the data release for this paper.

Due to the absence of evidence for coherent signals, we place upper limits from the posteriors on GW signal amplitudes as a function of τ for each search configuration. Figures 8 and 9 show upper limits $h_0^{95\%}$ for the case where $l = m = 2$, and $C_{21}^{95\%}$ and $C_{22}^{95\%}$ in the case for $l = 2, m = 1, 2$, respectively. In each figure, results are included from analyses assuming the rectangular time dependence from Equation (8) or the exponential dependence from Equation (9), and for unconstrained and constrained orientation angles, ι and ψ . All these upper limits were calculated using samples from the posterior distributions constrained within the τ bins indicated by the horizontal bars in each figure. The `CWInPy` line in Figure 5 corresponds to the exponential window and constrained orientation results, i.e., the dashed orange line in the first panel of Figure 8.

B. DETAILS OF THE TRANSIENT \mathcal{F} -STATISTIC ANALYSIS

Here we provide additional details of the setup and results of the transient \mathcal{F} -statistic search, as summarized in Section 5.3 of the main text. For O4, the search setup, distributed job management and post-processing from the previous O2 and O3 searches (Keitel et al. 2019; Abbott et al. 2022a; Modafferi et al. 2021, 2023) have been refactored as the `tCWip` (“transient continuous-wave investigation pipeline”) package (Keitel et al. 2025). This wraps CPU code in `LALSuite` (LVK 2025) for rectangular-window (constant amplitude) signals and a `PyCUDA` GPU implementation for exponentially decaying signals in `PyFstat` (Keitel & Ashton 2018; Keitel et al. 2021).

The same grid of templates is used to cover the frequency and frequency derivatives (up to \ddot{f}_{gw}) ranges for both amplitude evolution options. The grid spacings are calculated to yield a maximum metric mismatch (Prix 2007; Wette et al. 2008) of 0.02 over the full T_{obs} . The search ranges in all relevant parameters summarized in Table 3.

To determine whether there are significant candidates from one of the two analyses, we use the empirical background estimation procedure (Tenorio et al. 2022a) implemented in the `distromax` package (Tenorio et al. 2022b): We divide the search results into batches of 5000 templates, fit the distribution of the per-batch maximum detection statistics, and extrapolate this to obtain an estimate of the expected loudest outlier from the full template bank. The automated notching also described in Tenorio et al. (2022a) is used to reduce the impact of moderate-strength noise disturbances, with one iteration step for the rectangular-window search results and two steps for the exponential-window search results. For rectangular signals, the threshold is set at

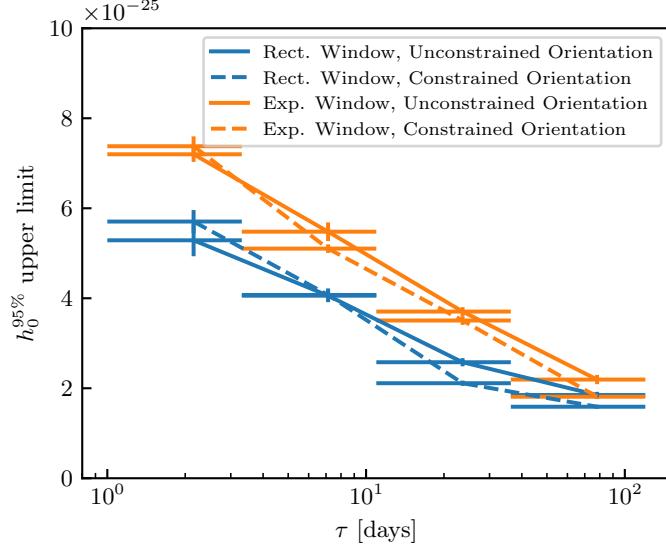


Figure 8. Upper limits from the fully-coherent targeted time-domain Bayesian search (CWInPy) in terms of $h_0^{95\%}$ as a function of signal duration τ . This search assumes emission only from the $l = m = 2$ mass quadrupole mode. “Rect. Window” lines correspond to the time dependence given in Equation (8), and “Exp. Window” lines to that in Equation (9). The solid and dashed lines indicate searches with unconstrained and constrained orientation angles, ι and ψ , respectively.

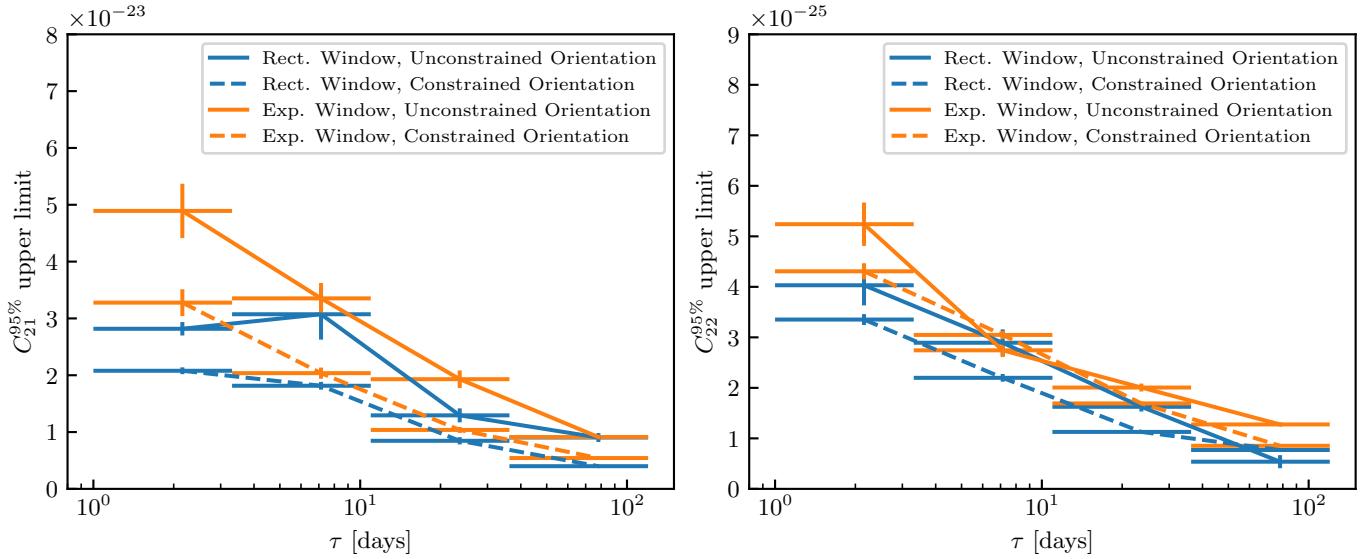


Figure 9. Upper limits from the time-domain Bayesian search (CWInPy) in terms of $C_{21}^{95\%}$ and $C_{22}^{95\%}$ as a function of signal duration, τ . This search assumes emission from the $l = 2, m = 1, 2$ mass quadrupole modes, where the first panel contains searches for $C_{21}^{95\%}$ and the second panel for $C_{22}^{95\%}$. “Rect. Window” lines correspond to the time dependence given in Equation (8), and “Exp. Window” lines to that in Equation (9). The solid and dashed lines indicate searches with unconstrained and constrained orientation angles, ι and ψ , respectively.

$\log_{10} \mathcal{B}_{tS/G}^{\text{thr}} \approx 10.02$ while the loudest candidate has $\max \log_{10} \mathcal{B}_{tS/G} \approx 9.65$. For the exponential-window search, the threshold is set to $\log_{10} \mathcal{B}_{tS/G}^{\text{thr}} \approx 9.30$ and the loudest candidate has $\max \log_{10} \mathcal{B}_{tS/G} \approx 9.20$.

We then obtain Frequentist 95% upper limits on the initial strain amplitude h_0 by simulating signals with rectangular windows or single exponential decay windows, following Equations (8) or (9). That is, these upper limits are not making statements about a specific model following the three-component decay observed in the radio, but about arbitrary CW-like transients within the search range as long as they fall within the range spanned by frequency and amplitude evolution templates. In fact, the effects of different amplitude evolutions on the sensitivity are relatively minor, see (Pri et al. 2011; Keitel et al. 2019), but still notable enough to motivate the use of two separate search configurations for rectangular and exponential window functions.

Table 3. Setup for the transient \mathcal{F} -statistic search.

parameter	value
T_{ref} [s]	1398372935
$t_{0,\text{min}}$ [s]	1398372935
Δt_0 [d]	2
τ_{min} [s]	3600
$\Delta\tau$ [d]	120
f_{gw} [Hz]	22.35453888321711
Δf_{gw} [Hz]	0.022365721744089156
\dot{f}_{gw} [Hz s $^{-1}$]	$-3.12874187125924 \times 10^{-11}$
$\Delta \dot{f}_{\text{gw}}$ [Hz s $^{-1}$]	$2.061675716980504 \times 10^{-13}$
\ddot{f}_{gw} [Hz s $^{-2}$]	$1.213100115859692 \times 10^{-21}$
$\Delta \ddot{f}_{\text{gw}}$ [Hz s $^{-2}$]	$2.639553012031674 \times 10^{-23}$
\dddot{f}_{gw} [Hz s $^{-3}$]	$2.5365827014929 \times 10^{-28}$
$\Delta \dddot{f}_{\text{gw}}$ [Hz s $^{-3}$]	$5.940719999999962 \times 10^{-30}$
N_{λ}	757016

NOTE— The reference time T_{ref} for the GW frequency and spin-down parameters f_{gw} , \dot{f}_{gw} , etc., is the first signal start time considered $t_{0,\text{min}}$, which in turn is set to the earliest SFT data timestamp in $T_{\text{gl}} \pm \min(\Delta T_{\text{gl}}, 1 \text{ d})$. The minimum signal duration τ_{min} is, for implementation reasons (PrixF 2015), set to two SFT lengths. Δt_0 and $\Delta\tau$ are the ranges of start times and durations considered. The minimum frequency f_{gw} and frequency bandwidth Δf_{gw} , and similarly for the three frequency derivatives, are set from the pulsar ephemeris (Table 1), using at each derivative order the maximum of (i) 0.001 times the central value of the parameter, (ii) 3σ uncertainties on it, (iii) for f_{gw} and \dot{f}_{gw} : the glitch step size. See Modafferi et al. (2021) for details on this setup procedure. N_{λ} is the total number of frequency-evolution templates.

For each window, the simulated signals cover a set of discrete τ values to sample the h_0 upper limits curves, as shown in Figure 10. At each τ value, a discrete set of h_0 values is covered to provide a detection efficiency curve, with 50 injections at each h_0 step. For the injections, the start time t_0 is randomly drawn from the ± 1 day search range, the frequency evolution parameters are drawn uniformly from the same ranges as covered by the search, and orientation angles are also uniformly sampled. For each h_0 , we then count as detectable those injections that exceed the threshold from the search with the corresponding window.

A sigmoid is fitted to the efficiency curve with `scipy` (Virtanen et al. 2020) and the $h_0^{95\%}$ value is inferred along with its error bar. As the detection statistic is a Bayes factor $\mathcal{B}_{ts/G}$ marginalized over start times and durations (PrixF et al. 2011), the upper limits at each injected τ value automatically include the full trials factor over the full allowed τ range of the search. (I.e., the search is not optimized differently depending on τ .)

C. DETAILS OF THE WPM ANALYSIS

For the WPM search summarized in Section 5.4, we provide here additional details on the setup and results, as well as a summary of investigations into the three outlier candidates produced by the search.

C.1. Details on WPM setup and results

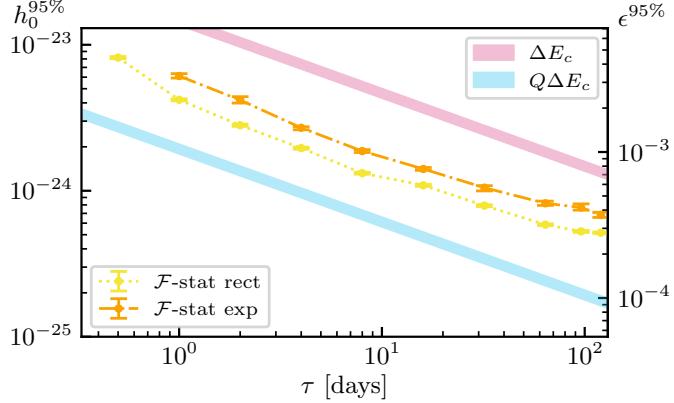


Figure 10. Upper limits from the transient \mathcal{F} -statistic search in terms of strain amplitude $h_0^{95\%}$ (left-hand axis) and NS ellipticity $\epsilon^{95\%}$ (right-hand axis), as a function of the signal duration parameter τ . Results are included both for signals with constant amplitude (“rect”), as per Equation (8), and for exponentially decaying signals (“exp”), as per Equation (9). The same indirect energy limits are shown for comparison as in Figure 5, corresponding to Equation (17) with $Q = 1$ (magenta) and $Q \approx 0.017$ (sky blue).

Table 4. Main parameters of the WPM search.

parameter	value
T_{start} [MJD]	60429.8466
T_{tot} [days]	101.7
T_{ref} [MJD]	60429.86975
$[f_{\text{start}}, f_{\text{stop}}]$ [Hz]	[22.306, 22.506] Hz
f_{gw} [Hz]	22.365772235265929
\dot{f}_{gw} [Hz s ⁻¹]	$-3.128837145670767 \times 10^{-11}$
\ddot{f}_{gw} [Hz s ⁻²]	$1.248409284183266 \times 10^{-21}$
\dddot{f}_{gw} [Hz s ⁻³]	$2.566286301492900 \times 10^{-28}$

NOTE— T_{start} is the time of the first sample of the analyzed dataset, 2000 seconds before the glitch. T_{ref} is the reference time for Vela’s position and rotational parameters, and corresponds to the glitch time. $[f_{\text{start}}, f_{\text{stop}}]$ defines the frequency interval over which the search is run. The values of the frequency f_{gw} and the first order spin-down, as used for heterodyning, include the variation due to the glitch, $2\Delta f_{\text{rot}}$, $2\Delta \dot{f}_{\text{rot}}$ respectively, see Table 1.

The WPM search uses the time and frequency parameters reported in Table 4. The segment duration is chosen so that the residual, uncorrected, spin-down due to a NS deformation induced by the glitch would produce a frequency variation confined within a single frequency bin. Moreover, the Doppler modulation correction is robust with respect to a frequency mismatch (i.e. a difference between the frequency used in the heterodyne step, f_{gw} in Table 4, and the actual signal frequency) as large as $\simeq 0.2$ Hz, larger than half the analyzed band.

A relevant issue in the search setup concerns the choice of the optimal observation window and segment duration. As a general rule, it is not convenient to consider an observation window which extends to times at which the signal amplitude is significantly decreased, as this would correspond to diluting the signal peaks in the noise.

We have studied the optimal window duration (considering a range from 0.5 days to T_{tot} , which is the total amount of data at our disposal) and segment duration (among four possible values of T_{seg} equal to 48000, 68000, 96000 and 172332 seconds) assuming both a signal with a single exponential decay, with various decay times, and a generalization of the model described

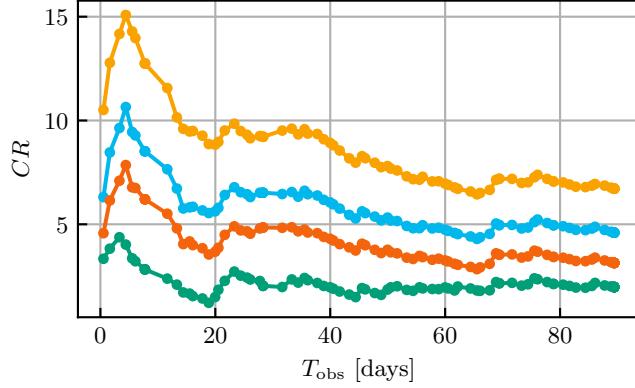


Figure 11. CR as a function of the observing window for the multi- τ signal model and using the optimal segment duration $T_{\text{seg}}^{\text{best}} = 96000$ seconds. The plot has been obtained injecting simulated signals into L1 data with four different signal amplitudes (going from the bottom to the top curve: $[2.1, 4.2, 6.3, 8.4] \times 10^{-24}$). The “optimal” observing window corresponds to the maximum of the CR , which is nearly independent of the signal amplitude, taking place at $T_{\text{obs}}^{\text{best}} \simeq 5.5$ days.

Table 5. Summary of the setup choices for the WPM search and 95% C.L. upper limits obtained for each search configuration.

τ [d]	$T_{\text{obs}}^{\text{best}}$ [d]	$T_{\text{seg}}^{\text{best}}$ [s]	$h_0^{\text{95\%,restr}}$	$h_0^{\text{95\%,unif}}$
15.11, 2.46, 0.39	5.5	96000	4.7×10^{-24}	4.5×10^{-24}
0.39	1.1	68000	4.5×10^{-24}	4.2×10^{-24}
2.46	5.0	96000	1.9×10^{-24}	2.0×10^{-24}
8	6.0	96000	1.3×10^{-24}	1.5×10^{-24}
15.11	5.5	96000	1.1×10^{-24}	1.1×10^{-24}
30	25	172332	1.1×10^{-24}	1.0×10^{-24}
1000	101.7	172332	5.6×10^{-25}	5.7×10^{-25}

NOTE— The first row corresponds to a search for a signal with amplitude varying in time according to a combination of three exponentials, with decay times inferred from the frequency evolution seen in radio observations. The other rows assume a signal amplitude described by a single exponential decay. The last row describes a signal with a nearly constant amplitude. $T_{\text{obs}}^{\text{best}}$ and $T_{\text{seg}}^{\text{best}}$ are, respectively, the optimal observation window and segment duration for the assumed signal model. These values depend not only on the signal model but also on the presence and position of gaps in the data. The penultimate column gives the upper limit for restricted priors on polarization parameters, and the last column for unrestricted uniform priors.

by Yim & Jones (2023), describing the amplitude evolution by a combination of three decaying exponentials, with decay times derived from the frequency evolution observed in the radio. Given one of these signal models, for each value of the observing window and of the segment duration, we have injected simulated signals of different amplitudes into O4b data and run the analyses, computing for each case the maximum CR . For each signal model, the values $T_{\text{obs}}^{\text{best}}$ and $T_{\text{seg}}^{\text{best}}$ that maximize the CR are our optimal choice.

As an example, in Figure 11 we show the CR as a function of the observing window for the multi- τ signal model, and for which the optimal segment duration is $T_{\text{seg}}^{\text{best}} = 96000$ s, for different signal amplitudes ($[2.1, 4.2, 6.3, 8.4] \times 10^{-24}$) injected into L1 O4b data. The maximum of the CR is at $T_{\text{obs}}^{\text{best}} \simeq 5.5$ days.

The full results of this optimization process are shown in columns 2 and 3 of Table 5, together with the numerical values of the final upper limits derived from the search. The full set of upper limits is also plotted in Figure 12.

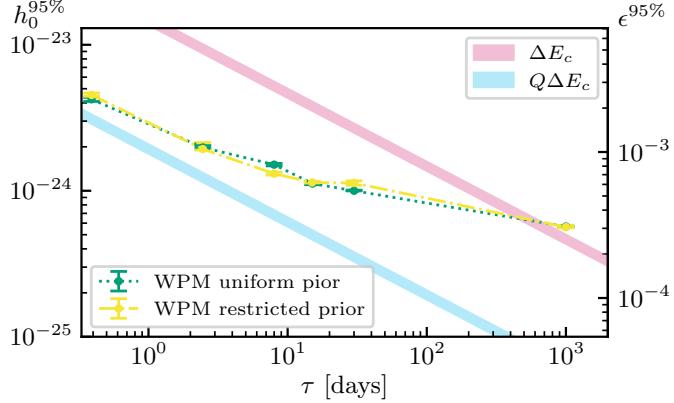


Figure 12. Upper limits (95% C.L.) from the WPM search in terms of strain amplitude and NS ellipticity immediately after the glitch, as a function of the signal decay time τ . The two lines correspond to the upper limits for, respectively, uniform and constrained priors on polarization parameters. See Section 5.2 for a discussion on the constrained parameters. For each τ the search is run on a time window $T_{\text{obs}}^{\text{best}}$ and using segments of duration $T_{\text{seg}}^{\text{best}}$ as shown in Table 5.

Upper limits are computed by simulated signal injections into the data, taking the amplitude such that 95% of the injected signals have a CR larger than the highest CR found in the analysis with the same setup, taking into account the specific 0.02 Hz sub-band in which each injection has been done.

C.2. Follow-up of outliers

In this section, we discuss the follow-up of the three outliers found in the WPM search.

The second and third outliers are likely due to an instrumental noise line. Indeed, by re-running the analysis *without* Doppler correction, their frequency shifts to 22.5 Hz and the corresponding CR increases significantly (by almost a factor of two), as expected for a detector disturbance.

The first outlier, which is more difficult to discard, has frequency 22.45538 Hz, spin-down $-3.108294 \times 10^{-11}$ Hz/s (at reference GPS time 1396742418), $CR = 6.65$ and has been found in the search configuration with $\tau = 8$ days, $T_{\text{obs}} = 6$ days and $T_{\text{seg}} = 96000$ seconds. Most of the follow-up tests we have done have not shown a clear incompatibility with an astrophysical signal.

As a first veto, we have again re-run the analysis switching off the Doppler modulation correction. This results in a decrease in CR from 6.65 to -0.96, which is compatible with what we expect for a real signal.

A second test consisted in running a semi-coherent search, using the semi-coherent 5-vector method (D’Antonio et al. 2023), with segment duration of 2, 3 and 6 days. Actually, these analyses produce three nearby candidates, with very similar CR , compatible with the original one, which are all considered in the next steps. The candidates remain significant when a segment duration of 2 days is used. For an astrophysical signal we expect the CR to increase with longer segment duration. This does *not* happen when passing from 2 to 3 days (the CR remains nearly unchanged), while it happens when 6 days are used (going from ~ 7.5 at 2 days to ~ 10.7 at 6 days).

We then repeated the analysis with a single data segment of 6 days shifting the starting time both backward and forward (-3.0, -1.5, -0.5, +0.5, +1.5, +3 days with respect to the signal reference starting time). The resulting CR s have been compared with what we would expect for a real signal which starts around the glitch time, both assuming an exponential decay with $\tau = 8$ days over a time window of 6 days and a constant amplitude signal of duration 5 or 7 days. We find that the CR steadily decreases both when the observing window is shifted backward and forward, becoming fully compatible with noise for a shift of ± 3 days. Overall, this test seems to indicate that there is a feature in the data which switches on around the glitch time and switches off after 5 to 7 days.

Finally, we have used the fully-coherent 5-vector method (Astone et al. 2010) to estimate the parameters of the potential signal associated to the candidates from the semi-coherent 5-vector step, using the frequency and spin-down values found in the $T_{\text{seg}} = 6$ days analysis. Estimated parameters for the three candidates we have analyzed are shown in Table 6. The last column, in particular, is an estimation of the so-called 5-vector “coherence”, a measure (between 0 and 1) of how much the data resembles

Table 6. Main parameters of the WPM outliers after coherent 5-vector follow-up.

Outlier	frequency [Hz]	spin-down [Hz/s]	h_0	η	ψ [deg]	CR	\mathcal{C}
1	22.455332	-3.316×10^{-11}	5.62×10^{-25}	-0.455	9.79	10.6	0.81
2	22.455333	-3.129×10^{-11}	7.72×10^{-25}	0.037	6.78	10.7	0.91
3	22.455334	-3.129×10^{-11}	4.64×10^{-25}	-0.20	9.80	10.7	0.54

NOTE—Parameter estimation, with the fully-coherent 5-vector method, for the three candidates associated to the first outlier from the main search. Rotational parameters are referred to MJD time 60429.8466 (GPS 1398457164.24).

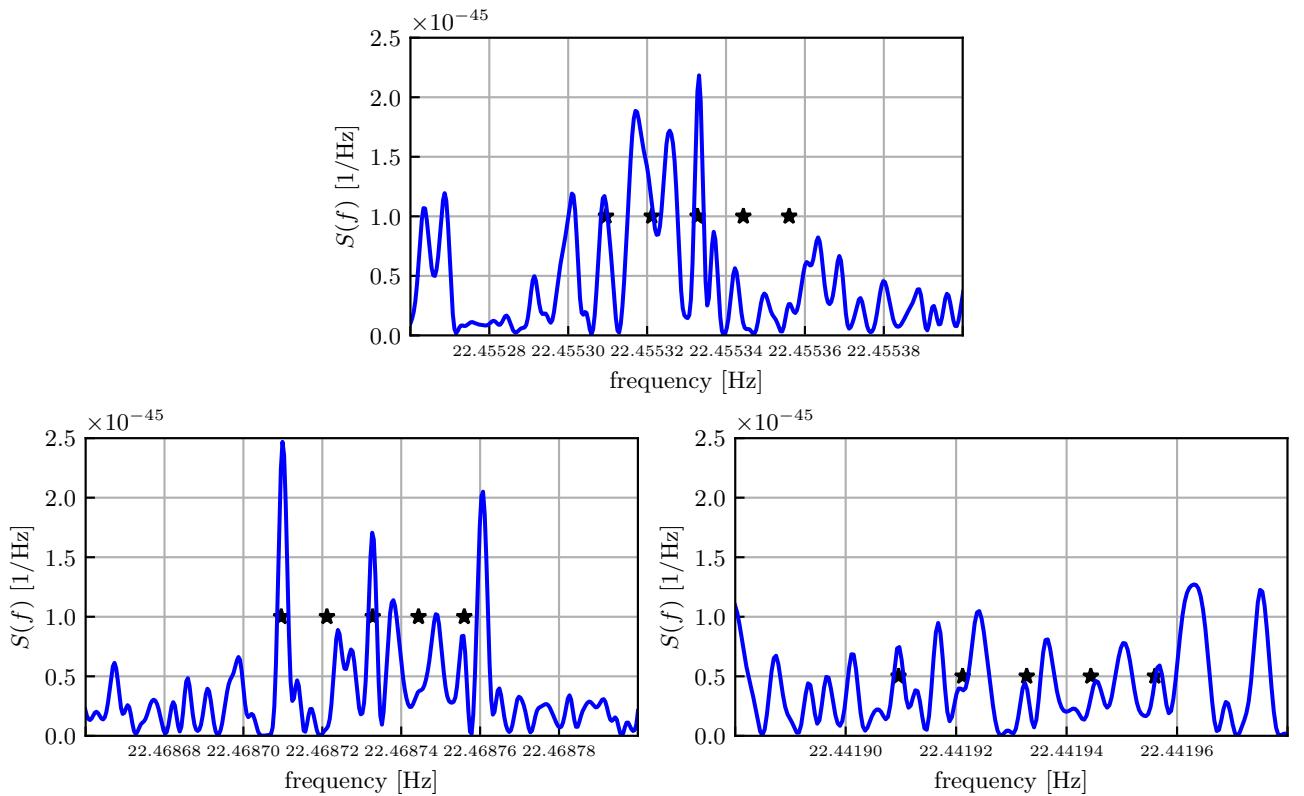


Figure 13. Top plot: power spectrum $S(f)$ of the LIGO data, after correction of the Doppler effect and spin-down of the second candidate in Table 6. The two other plots show the corrected power spectrum after the injection of a simulated signal with the same parameters as the second candidate in Table 6 but with frequency shifted by $+13.4$ mHz (bottom left) and -13.4 mHz (bottom right). The asterisks denote the position of the five expected peaks produced by the sidereal modulation.

the best matching template. The highest value we get, $\mathcal{C} = 0.91$ for the second candidate, corresponds to a false-alarm probability (in Gaussian noise) of about 0.012, after taking the trial factor into account.

We have then inspected the power spectrum of the data, after Doppler and spin-down correction: in presence of a sufficiently strong astrophysical signal, we expect to see the typical 5-vector signature, due to the signal's sidereal modulation. The spectrum after correcting the data for the parameters of the second candidate in Table 6 (that with the highest *coherence*) is shown in the top plot of Figure 13. The five asterisks denote the expected frequency of the five peaks (not the amplitude). We can see that at least two, out of five, peaks have a rather clear counterpart in the data. As a cross-check, we have computed the spectrum after injecting a simulated signal with the same parameters as those estimated for the second candidate, except for the frequency that has been shifted by some random amount. As an example in Figure 13 we show the spectra for shifts of ± 13.4 mHz. At the higher frequency, the spectrum looks slightly better when compared to the expected peak position, with two peaks clearly

Table 7. Search setup for the HMM search.

parameter	value
T_{start} [GPS]	1398458994
T_{ref} [GPS]	1398459095.043
T_{obs} [days]	90
T_{coh} [days]	0.375, 2.4374, 15.1042
f_{rot} [Hz]	11.1828595
\dot{f}_{rot} [Hz/s]	$-1.5540586 \times 10^{-11}$
\ddot{f}_{rot} [Hz/s ²]	6.242×10^{-22}
\dddot{f}_{rot} [Hz/s ³]	1.28×10^{-28}
Δf_{rot} [Hz]	0.05
N_T	240, 36, 5

NOTE— The GPS start time is rounded to the last SFT timestamp before the glitch arrival time at the detectors. The coherence times are similarly rounded down to the nearest multiple of 1800 s. The reference time is the glitch arrival time at the solar system barycenter, and the spin frequencies are extrapolated to this time. The frequency, bandwidth, and frequency derivative correspond to the Vela spin frequency; the values at twice the spin frequency are found by multiplying these quantities by two.

corresponding to features in the data (the first and the third peak, as for the real analysis candidate) and, possibly, also the fourth. At the lower frequency, on the other hand, the situation is much less clear, with only one convincing correspondence among the expected peak frequencies and the features in the data.

Overall, our tests for this outlier have been inconclusive. Our conclusion is that, even if a GW signal was really present in the data, it is very marginal, and a detection cannot be claimed with sufficient confidence.

D. DETAILS OF THE HMM ANALYSIS

Here we provide additional details concerning the HMM search. The template parameters, search band, and coherence times are summarized in Table 7.

As discussed in Section 5.5, empirical detection thresholds are set using one of two methods. In the first method, we perform the search on 10^4 realizations of synthetic Gaussian noise data to accumulate log-likelihood samples $\bar{\mathcal{L}}_i$, keeping the same search parameters given in Table 7. We set thresholds based on the Viterbi log-likelihood $\bar{\mathcal{L}}_{\text{th}}$ at which a desired fraction α of the samples derived from Gaussian noise simulations are above the threshold, $\bar{\mathcal{L}}_i > \bar{\mathcal{L}}_{\text{th}}$. In the second method, we set thresholds in exactly the same way, but instead the log-likelihood samples are derived from searching real detector data at 10^4 off-target sky positions. We sample off-target sky positions by shifting the search from the true Vela position by a random angle between $\pm(5^\circ\text{--}15^\circ)$ in RA and $\pm 5^\circ$ in Dec. In both methods, we set the false-alarm probability at $\alpha = 0.01/3$ for each choice of T_{coh} , such that the overall rate of false alarms, accounting for the trials factor from three different coherence times, is 1%. The resulting thresholds for each choice of T_{coh} and sampling method are shown in Table 8, where we also list the maximum statistics obtained from the single and dual harmonic HMM searches.

In simulating the injections used to set upper limits, the relative amplitudes of each harmonic depend on the angle, θ , between the NS axis of rotation and the principal moment of inertia, where $\theta = \pi/2$ corresponds to a GW emission at $f_{\text{gw}} = 2f_{\text{rot}}$ only

Table 8. Detection statistic thresholds for the HMM search.

Statistic	$T_{coh} = 0.375$ days	2.4374 days	15.1042 days
$\bar{\mathcal{L}}_{th,G}^{2F}$	7.05	9.33	17.14
$\bar{\mathcal{L}}_{th,OT}^{2F}$	6.99	9.32	17.06
$\bar{\mathcal{L}}_{max}^{2F}$	6.85	8.43	14.44
$\bar{\mathcal{L}}_{th,G}^{1F+2F}$	11.86	14.77	24.43
$\bar{\mathcal{L}}_{th,OT}^{1F+2F}$	11.5	14.38	24.09
$\bar{\mathcal{L}}_{max}^{1F+2F}$	11.46	13.56	19.59

NOTE—All log-likelihood thresholds are quoted at the level of 1% false-alarm probability, including trials factors.

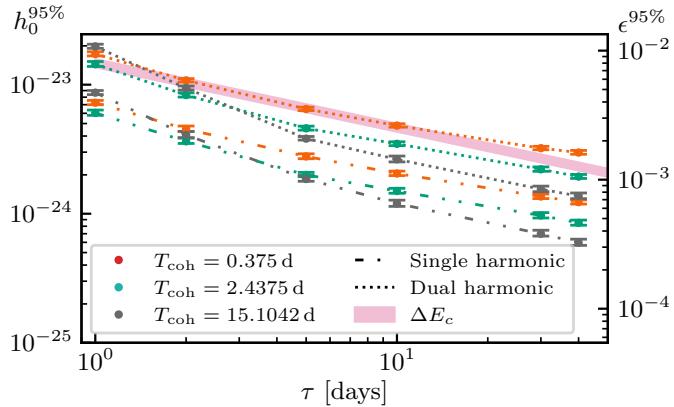


Figure 14. Upper limits on long-transient GWs from Vela inferred by the single (dash-dotted lines) and dual-harmonic (dotted lines) HMM analyses as a function of the exponential signal decay time parameter τ . Upper limits are expressed in terms of strain amplitude $h_0^{95\%}$ (left axis) and NS ellipticity $\epsilon^{95\%}$ (right axis). The magenta shaded comparison band labeled ΔE_c corresponds to Equation (17) with $Q = 1$.

(e.g., Jaranowski et al. 1998). We therefore fix $\theta = \pi/2$ for the single harmonic injections and $\theta = \pi/4$ for the dual harmonic injections, where the latter corresponds to roughly equal strain amplitudes at f_{rot} and $2f_{\text{rot}}$. Figure 14 shows upper limits on the strain amplitude and ellipticity, respectively, inferred by the HMM analysis. Different curves correspond to different choices of the coherence time and whether the signal model is assumed to have one component at twice the Vela spin frequency or at both once and twice the spin frequency. Unfortunately, because of the large disparity in the detector noise floors between 11 Hz and 22 Hz, the band at $f_{\text{gw}} = f_{\text{rot}}$ mostly contributes noise for any reasonable range of h_0 . Hence, our upper limits that assume dual harmonic emission are consistently worse by a factor of ~ 2.3 .

E. DETAILS OF THE JOINT NS PROPERTIES INFERENCE

To sample the parameter space for joint inference under the assumption of the Vela pulsar emitting GWs both from an f-mode and a transient mountain, we use the following priors. The NS mass prior is uniform in $[0.9 M_{\odot}, 2.6 M_{\odot}]$. The NS radius prior is uniform in [8 km, 16 km]. Constraints are enforced such that the mass and radius obey causality and the NS does not collapse to a black hole (see black contour in Figure 6). The prior for the glitch-induced spin change $\Delta f_{\text{rot}}/f_{\text{rot}}$ is a Gaussian distribution centered at 2.377×10^{-6} with $\sigma = 4.47 \times 10^{-11}$. We choose the inclination angle ι prior to be a double-peaked Gaussian distribution with peaks at 1.1048 and 2.0368 radians and $\sigma = 10.5$ mrad in order to match the astrophysical prior used in Section 5.2 and based on Ng & Romani (2008). The distance d prior is a Gaussian centered on 287 pc with $\sigma = 9.5$ pc, matching Dodson et al. (2003). The mountain decay time τ_{CW} prior is uniform from 1 hour to 120 days, matching the search durations covered in Section 5. The prior for the mountain energy scaling factor F_{CW} is log-uniform from 10^{-4} to 1.0 (none of the searches would be sensitive below the cutoff) and the prior for the f-mode energy scaling factor F_{fmode} is uniform between 0 and 1 (Figure 6) or fixed to 0.8 (Figure 7). For the model-specific results reported in Figure 7, we also set the prior on the

duration parameter τ_{CW} of CW-like signals to a Gaussian distribution peaked at twice the first decay time $\tau_d^1 = 15.1$ days of the radio frequency evolution, with $\sigma = 0.2$ days, to match the [Yim & Jones \(2020\)](#) model and Table 1.

We then combine the non-detection results from one representative pipeline each for the burst-type search for f-modes and the CW-like search for long transients from post-glitch mountain formation. For the f-mode, we use the ringdown injection curves generated by X-pipeline for different frequencies and damping times. These curves are linearly interpolated and subtracted from unity to give the probability that a waveform generated by a particular set of parameters would be detected by X-pipeline. For the mountain model, we use the posterior distribution on amplitude h_0 and CW damping time τ_{CW} produced by CWInPy for an exponential decay profile and constrained orientation, as described in appendix A. This likelihood distribution is interpreted as the probability that a given parameter would generate a waveform consistent with background noise.

For each sample, we compute the probability of non-detection using the X-pipeline-informed likelihood function and the probability of consistency with background using the CWInPy-informed likelihood function. We then re-weigh the prior samples with the product of these likelihoods according to Bayes' Theorem to give a posterior distribution on the sampled parameters consistent with search results. This re-weighting is done via rejection sampling, where the product of the two likelihoods is used as the probability that a given sample will not be rejected. Samples with higher likelihood are thus more likely to be preserved, and samples with lower likelihoods are more likely to be rejected.

Further details of the implementation are described in [Ball & Frey \(2025\)](#).

REFERENCES

Aasi, J., Abbott, B. P., Abbott, R., et al. 2015, CQGra, 32, 074001, doi: [10.1088/0264-9381/32/7/074001](https://doi.org/10.1088/0264-9381/32/7/074001)

Abac, A. G., Abouelfettouh, I., Acernese, F., et al. 2025a, arXiv e-prints, arXiv:2507.12282, doi: [10.48550/arXiv.2507.12282](https://doi.org/10.48550/arXiv.2507.12282)

—. 2025b, PhRvD, 112, 102005, doi: [10.1103/PhysRevD.112.102005](https://doi.org/10.1103/PhysRevD.112.102005)

—. 2025c, arXiv e-prints, arXiv:2508.18079, doi: [10.48550/arXiv.2508.18079](https://doi.org/10.48550/arXiv.2508.18079)

Abac, A. G., Abbott, R., Abouelfettouh, I., et al. 2025d, Astrophys. J., 985, 183, doi: [10.3847/1538-4357/adc681](https://doi.org/10.3847/1538-4357/adc681)

—. 2025e, ApJ, 983, 99, doi: [10.3847/1538-4357/adb3a0](https://doi.org/10.3847/1538-4357/adb3a0)

Abadie, J., Abbott, B. P., Abbott, R., et al. 2011, PhRvD, 83, 042001, doi: [10.1103/PhysRevD.83.042001](https://doi.org/10.1103/PhysRevD.83.042001)

Abbott, B., Abbott, R., Abbott, T., et al. 2017, PhRvL, 119, doi: [10.1103/PhysRevLett.119.161101](https://doi.org/10.1103/PhysRevLett.119.161101)

—. 2018, PhRvL, 121, doi: [10.1103/PhysRevLett.121.161101](https://doi.org/10.1103/PhysRevLett.121.161101)

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, ApJ, 839, 12, doi: [10.3847/1538-4357/aa677f](https://doi.org/10.3847/1538-4357/aa677f)

—. 2019, ApJ, 879, 10, doi: [10.3847/1538-4357/ab20cb](https://doi.org/10.3847/1538-4357/ab20cb)

—. 2020, Living Rev. Rel., 23, 3, doi: [10.1007/s41114-020-00026-9](https://doi.org/10.1007/s41114-020-00026-9)

Abbott, R., Abbott, T. D., Acernese, F., et al. 2021a, PhRvD, 104, 122004, doi: [10.1103/PhysRevD.104.122004](https://doi.org/10.1103/PhysRevD.104.122004)

—. 2021b, PhRvD, 104, 102001, doi: [10.1103/PhysRevD.104.102001](https://doi.org/10.1103/PhysRevD.104.102001)

—. 2022a, ApJ, 932, 133, doi: [10.3847/1538-4357/ac6ad0](https://doi.org/10.3847/1538-4357/ac6ad0)

Abbott, R., Abe, H., Acernese, F., et al. 2022b, ApJ, 935, 1, doi: [10.3847/1538-4357/ac6acf](https://doi.org/10.3847/1538-4357/ac6acf)

Abbott, R., Abbott, T. D., Acernese, F., et al. 2023, PhRvX, 13, 041039, doi: [10.1103/PhysRevX.13.041039](https://doi.org/10.1103/PhysRevX.13.041039)

Abbott, R., Abe, H., Acernese, F., et al. 2024, ApJ, 966, 137, doi: [10.3847/1538-4357/ad27d3](https://doi.org/10.3847/1538-4357/ad27d3)

Acernese, F., Agathos, M., Agatsuma, K., et al. 2015, CQGra, 32, 024001, doi: [10.1088/0264-9381/32/2/024001](https://doi.org/10.1088/0264-9381/32/2/024001)

Akmal, A., Pandharipande, V. R., & Ravenhall, D. G. 1998, PhRvC, 58, 1804, doi: [10.1103/PhysRevC.58.1804](https://doi.org/10.1103/PhysRevC.58.1804)

Akutsu, T., Ando, M., Arai, K., et al. 2021, Progress of Theoretical and Experimental Physics, 2021, 05A101, doi: [10.1093/ptep/ptaa125](https://doi.org/10.1093/ptep/ptaa125)

Allen, B., Goetz, E., Keitel, D., et al. 2024, SFT Data Format Version 2–3 Specification, Tech. Rep. LIGO-T040164-v3, LIGO Laboratory. <https://dcc.ligo.org/T040164-v3/public>

Andersson, N., & Kokkotas, K. D. 1998, MNRAS, 299, 1059, doi: [10.1046/j.1365-8711.1998.01840.x](https://doi.org/10.1046/j.1365-8711.1998.01840.x)

Antonopoulou, D., Haskell, B., & Espinoza, C. M. 2022, Reports on Progress in Physics, 85, 126901, doi: [10.1088/1361-6633/ac9ced](https://doi.org/10.1088/1361-6633/ac9ced)

Ashton, G., Keitel, D., Prix, R., Tenorio, R., & Ferrer, M.-A. 2025, PyFstat, v2.2.1, Zenodo, doi: [10.5281/zenodo.14919610](https://doi.org/10.5281/zenodo.14919610)

Ashton, G., Hübner, M., Lasky, P. D., et al. 2019, ApJS, 241, 27, doi: [10.3847/1538-4365/ab06fc](https://doi.org/10.3847/1538-4365/ab06fc)

Astone, P., Colla, A., D'Antonio, S., Frasca, S., & Palomba, C. 2014, PhRvD, 90, 042002, doi: [10.1103/PhysRevD.90.042002](https://doi.org/10.1103/PhysRevD.90.042002)

Astone, P., D'Antonio, S., Frasca, S., & Palomba, C. 2010, CQGra, 27, 194016, doi: [10.1088/0264-9381/27/19/194016](https://doi.org/10.1088/0264-9381/27/19/194016)

Astone, P., Frasca, S., & Palomba, C. 2005, CQGra, 22, S1197, doi: [10.1088/0264-9381/22/18/S34](https://doi.org/10.1088/0264-9381/22/18/S34)

Baiko, D. A., & Chugunov, A. I. 2018, MNRAS, 480, 5511, doi: [10.1093/mnras/sty2259](https://doi.org/10.1093/mnras/sty2259)

Ball, M., & Frey, R. 2025, arXiv e-prints, arXiv:2508.15965, doi: [10.48550/arXiv.2508.15965](https://doi.org/10.48550/arXiv.2508.15965)

Barr, E. D., Dutta, A., Freire, P. C. C., et al. 2024, Science, 383, 275, doi: [10.1126/science.adg3005](https://doi.org/10.1126/science.adg3005)

Basu, A., Shaw, B., Antonopoulou, D., et al. 2022, MNRAS, 510, 4049, doi: [10.1093/mnras/stab3336](https://doi.org/10.1093/mnras/stab3336)

Behnke, B., Papa, M. A., & Prix, R. 2015, PhRvD, 91, 064007, doi: [10.1103/PhysRevD.91.064007](https://doi.org/10.1103/PhysRevD.91.064007)

Bennett, M. F., van Eysden, C. A., & Melatos, A. 2010, MNRAS, 409, 1705, doi: [10.1111/j.1365-2966.2010.17416.x](https://doi.org/10.1111/j.1365-2966.2010.17416.x)

Bransgrove, A., Beloborodov, A. M., & Levin, Y. 2020, ApJ, 897, 173, doi: [10.3847/1538-4357/ab93b7](https://doi.org/10.3847/1538-4357/ab93b7)

Brun, R., & Rademakers, F. 1997, Nucl. Instrum. Meth. A, 389, 81, doi: [10.1016/s0168-9002\(97\)00048-x](https://doi.org/10.1016/s0168-9002(97)00048-x)

Capote, E., Jia, W., Aritomi, N., et al. 2025, PhRvD, 111, 062002, doi: [10.1103/PhysRevD.111.062002](https://doi.org/10.1103/PhysRevD.111.062002)

Cheunchitra, T., Melatos, A., Carlin, J. B., & Howitt, G. 2024, MNRAS, 528, 1360, doi: [10.1093/mnras/stae130](https://doi.org/10.1093/mnras/stae130)

Cordes, J. M., Downs, G. S., & Krause-Polstorf, J. 1988, ApJ, 330, 847, doi: [10.1086/166518](https://doi.org/10.1086/166518)

Cutler, C., & Schutz, B. F. 2005, PhRvD, 72, 063006, doi: [10.1103/PhysRevD.72.063006](https://doi.org/10.1103/PhysRevD.72.063006)

D'Antonio, S., et al. 2023, PhRvD, 108, 122001, doi: [10.1103/PhysRevD.108.122001](https://doi.org/10.1103/PhysRevD.108.122001)

Davis, D., Neunzert, A., Goetz, E., et al. 2025, Self-gating of O4a h(t) for use in continuous-wave searches, Tech. Rep. LIGO-T2400003, LIGO Laboratory. <https://dcc.ligo.org/T2400003/public>

De Lillo, F., Suresh, J., Depasse, A., et al. 2023, PhRvD, 107, 102001, doi: [10.1103/PhysRevD.107.102001](https://doi.org/10.1103/PhysRevD.107.102001)

Dodson, R., Legge, D., Reynolds, J. E., & McCulloch, P. M. 2003, ApJ, 596, 1137, doi: [10.1086/378089](https://doi.org/10.1086/378089)

Dodson, R., Lewis, D., & McCulloch, P. 2007, Astrophysics and Space Science, 308, 585, doi: [10.1007/s10509-007-9372-4](https://doi.org/10.1007/s10509-007-9372-4)

Douchin, F., & Haensel, P. 2001, Astronomy & Astrophysics, 380, 151, doi: [10.1051/0004-6361:20011402](https://doi.org/10.1051/0004-6361:20011402)

Drago, M., et al. 2021, SoftwareX, 14, 100678, doi: [10.1016/j.softx.2021.100678](https://doi.org/10.1016/j.softx.2021.100678)

Dreissigacker, C., Prix, R., & Wette, K. 2018, PhRvD, 98, 084058, doi: [10.1103/PhysRevD.98.084058](https://doi.org/10.1103/PhysRevD.98.084058)

Dunn, L., Melatos, A., Espinoza, C. M., Antonopoulou, D., & Dodson, R. 2023, MNRAS, 522, 5469, doi: [10.1093/mnras/stad1335](https://doi.org/10.1093/mnras/stad1335)

Dupuis, R. J., & Woan, G. 2005, PhRvD, 72, 102002, doi: [10.1103/PhysRevD.72.102002](https://doi.org/10.1103/PhysRevD.72.102002)

Echeverria, F. 1989, PhRvD, 40, 3194, doi: [10.1103/PhysRevD.40.3194](https://doi.org/10.1103/PhysRevD.40.3194)

Edwards, R. T., Hobbs, G. B., & Manchester, R. N. 2006, MNRAS, 372, 1549, doi: [10.1111/j.1365-2966.2006.10870.x](https://doi.org/10.1111/j.1365-2966.2006.10870.x)

Espinoza, C. M., Antonopoulou, D., Dodson, R., Stepanova, M., & Scherer, A. 2021, A&A, 649, C2, doi: [10.1051/0004-6361/202039044e](https://doi.org/10.1051/0004-6361/202039044e)

Espinoza, C. M., Lyne, A. G., Stappers, B. W., & Kramer, M. 2011, MNRAS, 414, 1679, doi: [10.1111/j.1365-2966.2011.18503.x](https://doi.org/10.1111/j.1365-2966.2011.18503.x)

Finn, L. S. 1992, PhRvD, 46, 5236, doi: [10.1103/PhysRevD.46.5236](https://doi.org/10.1103/PhysRevD.46.5236)

Fuentes, J. R., Espinoza, C. M., Reisenegger, A., et al. 2017, A&A, 608, A131, doi: [10.1051/0004-6361/201731519](https://doi.org/10.1051/0004-6361/201731519)

Ganapathy, D., Jia, W., Nakano, M., et al. 2023, PhRvX, 13, 041021, doi: [10.1103/PhysRevX.13.041021](https://doi.org/10.1103/PhysRevX.13.041021)

Gancio, G., Lousto, C. O., Combi, L., et al. 2020, A&A, 633, A84, doi: [10.1051/0004-6361/201936525](https://doi.org/10.1051/0004-6361/201936525)

Glampedakis, K., & Gualtieri, L. 2018, in Astrophysics and Space Science Library, Vol. 457, Astrophysics and Space Science Library, ed. L. Rezzolla, P. Pizzochero, D. I. Jones, N. Rea, & I. Vidaña, 673, doi: [10.1007/978-3-319-97616-7-12](https://doi.org/10.1007/978-3-319-97616-7-12)

Goetz, E., & Riles, K. 2024, Segments used for creating standard SFTs in O4 data, Tech. Rep. LIGO-T2400058-v3, LIGO Laboratory. <https://dcc.ligo.org/T2400058-v3/public>

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Nature, 585, 357, doi: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2)

Haskell, B. 2018, in Pulsar Astrophysics the Next Fifty Years, ed. P. Weltevrede, B. B. P. Perera, L. L. Preston, & S. Sanidas, Vol. 337, 203–208, doi: [10.1017/S1743921317010663](https://doi.org/10.1017/S1743921317010663)

Haskell, B., & Jones, D. I. 2024, Astroparticle Physics, 157, 102921, doi: [10.1016/j.astropartphys.2023.102921](https://doi.org/10.1016/j.astropartphys.2023.102921)

Haskell, B., & Melatos, A. 2015, International Journal of Modern Physics D, 24, 1530008, doi: [10.1142/S0218271815300086](https://doi.org/10.1142/S0218271815300086)

Ho, W. C. G., Jones, D. I., Andersson, N., & Espinoza, C. M. 2020, PhRvD, 101, 103009, doi: [10.1103/PhysRevD.101.103009](https://doi.org/10.1103/PhysRevD.101.103009)

Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, MNRAS, 369, 655, doi: [10.1111/j.1365-2966.2006.10302.x](https://doi.org/10.1111/j.1365-2966.2006.10302.x)

Hotan, A. W., van Straten, W., & Manchester, R. N. 2004, PASA, 21, 302, doi: [10.1071/AS04022](https://doi.org/10.1071/AS04022)

Hunter, J. D. 2007, CSE, 9, 90, doi: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55)

Jaranowski, P., & Królak, A. 2010, CQGra, 27, 194015, doi: [10.1088/0264-9381/27/19/194015](https://doi.org/10.1088/0264-9381/27/19/194015)

Jaranowski, P., Królak, A., & Schutz, B. F. 1998, PhRvD, 58, 063001, doi: [10.1103/PhysRevD.58.063001](https://doi.org/10.1103/PhysRevD.58.063001)

Jia, W., Xu, V., Kuns, K., et al. 2024, Science, 385, 1318, doi: [10.1126/science.ado8069](https://doi.org/10.1126/science.ado8069)

Johnson-McDaniel, N. K., & Owen, B. J. 2013, PhRvD, 88, 044004, doi: [10.1103/PhysRevD.88.044004](https://doi.org/10.1103/PhysRevD.88.044004)

Jones, D. I. 2010, MNRAS, 402, 2503, doi: [10.1111/j.1365-2966.2009.16059.x](https://doi.org/10.1111/j.1365-2966.2009.16059.x)

Jones, D. I., & Andersson, N. 2001, MNRAS, 324, 811, doi: [10.1046/j.1365-8711.2001.04251.x](https://doi.org/10.1046/j.1365-8711.2001.04251.x)

—. 2002, MNRAS, 331, 203, doi: [10.1046/j.1365-8711.2002.05180.x](https://doi.org/10.1046/j.1365-8711.2002.05180.x)

Keer, L., & Jones, D. I. 2015, MNRAS, 446, 865, doi: [10.1093/mnras/stu2123](https://doi.org/10.1093/mnras/stu2123)

Keitel, D., & Ashton, G. 2018, CQGra, 35, 205003, doi: [10.1088/1361-6382/aade34](https://doi.org/10.1088/1361-6382/aade34)

Keitel, D., Singh, N., & Ferrer, M.-A. 2025, tCWip – transient continuous-wave investigation pipeline, v0.5.5, open source software, doi: [10.5281/zenodo.1798985](https://doi.org/10.5281/zenodo.1798985)

Keitel, D., Tenorio, R., Ashton, G., & Prix, R. 2021, The Journal of Open Source Software, 6, 3000, doi: [10.21105/joss.03000](https://doi.org/10.21105/joss.03000)

Keitel, D., Woan, G., Pitkin, M., et al. 2019, PhRvD, 100, 064058, doi: [10.1103/PhysRevD.100.064058](https://doi.org/10.1103/PhysRevD.100.064058)

Klimenko, S., et al. 2016, PhRvD, 93, 042004, doi: [10.1103/PhysRevD.93.042004](https://doi.org/10.1103/PhysRevD.93.042004)

Kokkotas, K. D., & Schmidt, B. G. 1999, Living Reviews in Relativity, 2, 2, doi: [10.12942/lrr-1999-2](https://doi.org/10.12942/lrr-1999-2)

Lackey, B. D., Nayyar, M., & Owen, B. J. 2006, PhRvD, 73, 024021, doi: [10.1103/PhysRevD.73.024021](https://doi.org/10.1103/PhysRevD.73.024021)

Large, M. I., Vaughan, A. E., & Mills, B. Y. 1968, Nature, 220, 340, doi: [10.1038/220340a0](https://doi.org/10.1038/220340a0)

Lattimer, J. M. 2021, Ann. Rev. Nucl. Part. Sci., 71, 433, doi: [10.1146/annurev-nucl-102419-124827](https://doi.org/10.1146/annurev-nucl-102419-124827)

Levin, Y., & van Hoven, M. 2011, MNRAS, 418, 659, doi: [10.1111/j.1365-2966.2011.19515.x](https://doi.org/10.1111/j.1365-2966.2011.19515.x)

Link, B., Epstein, R. I., & Lattimer, J. M. 2000, in Astrophysics and Space Science Library, Vol. 254, Stellar Astrophysics, ed. K. S. Cheng, H. F. Chau, K. L. Chan, & K. C. Leung, 117, doi: [10.1007/978-94-010-0878-5_14](https://doi.org/10.1007/978-94-010-0878-5_14)

Link, B., Epstein, R. I., & van Riper, K. A. 1992, Nature, 359, 616, doi: [10.1038/359616a0](https://doi.org/10.1038/359616a0)

Lopez, D., Tiwari, S., Drago, M., et al. 2022, PhRvD, 106, 103037, doi: [10.1103/PhysRevD.106.103037](https://doi.org/10.1103/PhysRevD.106.103037)

Lopez Armengol, F. G., Lousto, C. O., del Palacio, S., et al. 2019, The Astronomer's Telegram, 12482, 1, <https://www.astronomerstelegram.org/?read=12482>

Lousto, C. O., Missel, R., Zubietta, E., et al. 2024, in Revista Mexicana de Astronomia y Astrofisica Conference Series, Vol. 56, Revista Mexicana de Astronomia y Astrofisica Conference Series, 134–144, doi: [10.22201/ia.14052059p.2024.56.25](https://doi.org/10.22201/ia.14052059p.2024.56.25)

LVK. 2025, LVK Algorithm Library - LALSuite, Free software (GPL), doi: [10.7935/GT1W-FZ16](https://doi.org/10.7935/GT1W-FZ16)

Macleod, D. M., Areeda, J. S., Coughlin, S. B., Massinger, T. J., & Urban, A. L. 2021, SoftwareX, 13, 100657, doi: [10.1016/j.softx.2021.100657](https://doi.org/10.1016/j.softx.2021.100657)

Macquet, A., Bizouard, M. A., Christensen, N., & Coughlin, M. 2021, PhRvD, 104, 102005, doi: [10.1103/PhysRevD.104.102005](https://doi.org/10.1103/PhysRevD.104.102005)

Martini, A., et al. 2025, arXiv e-prints, <https://arxiv.org/abs/2510.21411>

Melatos, A., Douglass, J. A., & Simula, T. P. 2015, ApJ, 807, 132, doi: [10.1088/0004-637X/807/2/132](https://doi.org/10.1088/0004-637X/807/2/132)

Modafferri, L. M., Moragues, J., & Keitel, D. 2021, in J. Phys. Conf. Ser., Vol. 2156, J. Phys. Conf. Ser. (IOP), 012079, doi: [10.1088/1742-6596/2156/1/012079](https://doi.org/10.1088/1742-6596/2156/1/012079)

Modafferri, L. M., Tenorio, R., & Keitel, D. 2023, PhRvD, 108, 023005, doi: [10.1103/PhysRevD.108.023005](https://doi.org/10.1103/PhysRevD.108.023005)

Moragues, J., Modafferri, L. M., Tenorio, R., & Keitel, D. 2023, MNRAS, 519, 5161, doi: [10.1093/mnras/stac3665](https://doi.org/10.1093/mnras/stac3665)

Necula, V., Klimenko, S., & Mitselmakher, G. 2012, J. Phys. Conf. Ser., 363, 012032, doi: [10.1088/1742-6596/363/1/012032](https://doi.org/10.1088/1742-6596/363/1/012032)

Ng, C. Y., & Romani, R. W. 2008, ApJ, 673, 411, doi: [10.1086/523935](https://doi.org/10.1086/523935)

Palfreyman, J. 2024, The Astronomer's Telegram, 16615, 1, <https://www.astronomerstelegram.org/?read=16615>

Palfreyman, J., Dickey, J. M., Hotan, A., Ellingsen, S., & van Straten, W. 2018, Nature, 556, 219, doi: [10.1038/s41586-018-0001-x](https://doi.org/10.1038/s41586-018-0001-x)

Piccinni, O. J., Astone, P., D'Antonio, S., et al. 2019, CQGra, 36, 015008, doi: [10.1088/1361-6382/aaefb5](https://doi.org/10.1088/1361-6382/aaefb5)

Pitkin, M. 2022, The Journal of Open Source Software, 7, 4568, doi: [10.21105/joss.04568](https://doi.org/10.21105/joss.04568)

Pitkin, M., Isi, M., Veitch, J., & Woan, G. 2017, arXiv e-prints, arXiv:1705.08978, doi: [10.48550/arXiv.1705.08978](https://doi.org/10.48550/arXiv.1705.08978)

Pradhan, B. K., Chatterjee, D., Lanoye, M., & Jaikumar, P. 2022, PhRvC, 106, 015805, doi: [10.1103/PhysRevC.106.015805](https://doi.org/10.1103/PhysRevC.106.015805)

Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., et al. 2018, AJ, 156, 123, doi: [10.3847/1538-3881/aabc4f](https://doi.org/10.3847/1538-3881/aabc4f)

Prix, R. 2007, PhRvD, 75, 023004, doi: [10.1103/PhysRevD.75.023004](https://doi.org/10.1103/PhysRevD.75.023004)

Prix, R. 2015, The F-statistic and its implementation in ComputeFStatistic_v2, Tech. Rep. LIGO-T0900149, LIGO Laboratory, <https://dcc.ligo.org/T0900149/public>

Prix, R., Giampaolis, S., & Messenger, C. 2011, PhRvD, 84, 023007, doi: [10.1103/PhysRevD.84.023007](https://doi.org/10.1103/PhysRevD.84.023007)

Raaijmakers, G., Greif, S. K., Hebeler, K., et al. 2021, ApJL, 918, L29, doi: [10.3847/2041-8213/ac089a](https://doi.org/10.3847/2041-8213/ac089a)

Radhakrishnan, V., & Manchester, R. N. 1969, Nature, 222, 228, doi: [10.1038/222228a0](https://doi.org/10.1038/222228a0)

Ransom, S. 2011, PRESTO: PulsaR Exploration and Search Toolkit, Astrophysics Source Code Library, record ascl:1107.017, <https://ascl.net/1107.017>

Reichley, P. E., & Downs, G. S. 1969, Nature, 222, 229, doi: [10.1038/222229a0](https://doi.org/10.1038/222229a0)

Riles, K. 2023, Living Reviews in Relativity, 26, 3, doi: [10.1007/s41114-023-00044-3](https://doi.org/10.1007/s41114-023-00044-3)

Riley, T. E., Watts, A. L., Ray, P. S., et al. 2021, ApJL, 918, L27, doi: [10.3847/2041-8213/ac0a81](https://doi.org/10.3847/2041-8213/ac0a81)

Robitaille, T. P., Tollerud, E. J., Greenfield, P., et al. 2013, A&A, 558, A33, doi: [10.1051/0004-6361/201322068](https://doi.org/10.1051/0004-6361/201322068)

Singh, A. 2017, PhRvD, 95, 024022, doi: [10.1103/PhysRevD.95.024022](https://doi.org/10.1103/PhysRevD.95.024022)

Soni, S., Berger, B. K., Davis, D., et al. 2025, CQGra, 42, 085016, doi: [10.1088/1361-6382/adc4b6](https://doi.org/10.1088/1361-6382/adc4b6)

Sosa-Fiscella, V., Zubia, E., del Palacio, S., et al. 2021, The Astronomer's Telegram, 14806, 1.
<https://www.astronomerstelegram.org/?read=14806>

Speagle, J. S. 2020, MNRAS, 493, 3132, doi: [10.1093/mnras/staa278](https://doi.org/10.1093/mnras/staa278)

Sun, L., Melatos, A., & Lasky, P. D. 2019, PhRvD, 99, 123010, doi: [10.1103/PhysRevD.99.123010](https://doi.org/10.1103/PhysRevD.99.123010)

Sun, L., Melatos, A., Suvorova, S., Moran, W., & Evans, R. J. 2018, PhRvD, 97, 043013, doi: [10.1103/PhysRevD.97.043013](https://doi.org/10.1103/PhysRevD.97.043013)

Sun, L., et al. 2020, CQGra, 37, 225008, doi: [10.1088/1361-6382/abb14e](https://doi.org/10.1088/1361-6382/abb14e)

Sun, L., Goetz, E., Kissel, J. S., et al. 2021, arXiv e-prints, arXiv:2107.00129, doi: [10.48550/arXiv.2107.00129](https://doi.org/10.48550/arXiv.2107.00129)

Sutton, P. J., et al. 2010, New J. Phys., 12, 053034, doi: [10.1088/1367-2630/12/5/053034](https://doi.org/10.1088/1367-2630/12/5/053034)

Suvorova, S., Sun, L., Melatos, A., Moran, W., & Evans, R. J. 2016, PhRvD, 93, 123009, doi: [10.1103/PhysRevD.93.123009](https://doi.org/10.1103/PhysRevD.93.123009)

Taylor, J. H., & Weisberg, J. M. 1982, ApJ, 253, 908, doi: [10.1086/159690](https://doi.org/10.1086/159690)

Tenorio, R., Modafferi, L. M., Keitel, D., & Sintes, A. M. 2022a, PhRvD, 105, 044029, doi: [10.1103/PhysRevD.105.044029](https://doi.org/10.1103/PhysRevD.105.044029)

—. 2022b, distromax, v1.1.0, Zenodo, doi: [10.5281/zenodo.7331554](https://doi.org/10.5281/zenodo.7331554)

Vajente, G., Huang, Y., Isi, M., et al. 2020, PhRvD, 101, 042003, doi: [10.1103/PhysRevD.101.042003](https://doi.org/10.1103/PhysRevD.101.042003)

van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., et al. 2014, PeerJ, 2, e453, doi: [10.7717/peerj.453](https://doi.org/10.7717/peerj.453)

van Eysden, C. A., & Melatos, A. 2008, CQGra, 25, 225020, doi: [10.1088/0264-9381/25/22/225020](https://doi.org/10.1088/0264-9381/25/22/225020)

van Straten, W., & Bailes, M. 2011, PASA, 28, 1, doi: [10.1071/AS10021](https://doi.org/10.1071/AS10021)

Vargas, A. F., & Melatos, A. 2023, MNRAS, 522, 4880, doi: [10.1093/mnras/stad1301](https://doi.org/10.1093/mnras/stad1301)

Viets, A., & Wade, M. 2021, Subtracting Narrow-band Noise from LIGO Strain Data in the Third Observing Run, Tech. Rep. LIGO-T2100058-v5, LIGO Laboratory.
<https://dcc.ligo.org/T2100058-v5/public>

Viets, A., Wade, M., Urban, A., et al. 2018, CQGra, 35, 095015

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261, doi: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2)

Viterbi, A. 1967, IEEE Transactions on Information Theory, 13, 260, doi: [10.1109/TIT.1967.1054010](https://doi.org/10.1109/TIT.1967.1054010)

Wade, M., Betzwieser, J., Bhattacharjee, D., et al. 2025, CQGra, 42, 215016, doi: [10.1088/1361-6382/ae1095](https://doi.org/10.1088/1361-6382/ae1095)

Warszawski, L., & Melatos, A. 2012, MNRAS, 423, 2058, doi: [10.1111/j.1365-2966.2012.20977.x](https://doi.org/10.1111/j.1365-2966.2012.20977.x)

Was, M., Sutton, P. J., Jones, G., & Leonor, I. 2012, PhRvD, 86, 022003, doi: [10.1103/PhysRevD.86.022003](https://doi.org/10.1103/PhysRevD.86.022003)

Was, M., Bizouard, M.-A., Brisson, V., et al. 2010, CQGra, 27, 015005, doi: [10.1088/0264-9381/27/1/015005](https://doi.org/10.1088/0264-9381/27/1/015005)

Wette, K. 2020, SoftwareX, 12, 100634, doi: [10.1016/j.softx.2020.100634](https://doi.org/10.1016/j.softx.2020.100634)

—. 2023, Astropart. Phys., 153, 102880, doi: [10.1016/j.astropartphys.2023.102880](https://doi.org/10.1016/j.astropartphys.2023.102880)

Wette, K., Owen, B. J., Allen, B., et al. 2008, CQGra, 25, 235011, doi: [10.1088/0264-9381/25/23/235011](https://doi.org/10.1088/0264-9381/25/23/235011)

Wilson, O. H., & Ho, W. C. G. 2024, PhRvD, 109, 083006, doi: [10.1103/PhysRevD.109.083006](https://doi.org/10.1103/PhysRevD.109.083006)

Wiringa, R. B., Fiks, V., & Fabrocini, A. 1988, PhRvC, 38, 1010, doi: [10.1103/physrevc.38.1010](https://doi.org/10.1103/physrevc.38.1010)

Yagi, K., & Yunes, N. 2017, Physics Reports, 681, 1, doi: <https://doi.org/10.1016/j.physrep.2017.03.002>

Yim, G., & Jones, D. I. 2020, MNRAS, 498, 3138, doi: [10.1093/mnras/staa2534](https://doi.org/10.1093/mnras/staa2534)

—. 2023, MNRAS, 518, 4322, doi: [10.1093/mnras/stac3405](https://doi.org/10.1093/mnras/stac3405)

Yim, G., Shao, L., & Xu, R. 2024, MNRAS, 532, 3893, doi: [10.1093/mnras/stae1659](https://doi.org/10.1093/mnras/stae1659)

Yu, M., Manchester, R. N., Hobbs, G., et al. 2013, MNRAS, 429, 688, doi: [10.1093/mnras/sts366](https://doi.org/10.1093/mnras/sts366)

Zhou, S., Gügercinoğlu, E., Yuan, J., Ge, M., & Yu, C. 2022, Universe, 8, 641, doi: [10.3390/universe8120641](https://doi.org/10.3390/universe8120641)

Zimmermann, M., & Szedenits, E. 1979, PhRvD, 20, 351, doi: [10.1103/PhysRevD.20.351](https://doi.org/10.1103/PhysRevD.20.351)

Zubia, E., Furlan, S. B. A., Palacio, S. d., et al. 2024a, The Astronomer's Telegram, 16608, 1.
<https://www.astronomerstelegram.org/?read=16608>

Zubia, E., García, F., del Palacio, S., et al. 2024b, A&A, 689, A191, doi: [10.1051/0004-6361/202450441](https://doi.org/10.1051/0004-6361/202450441)

Zubia, E., Missel, R., Sosa Fiscella, V., et al. 2023, MNRAS, 521, 4504, doi: [10.1093/mnras/stad723](https://doi.org/10.1093/mnras/stad723)

Zubia, E., Missel, R., Araujo Furlan, S. B., et al. 2025, A&A, 698, A72, doi: [10.1051/0004-6361/202554098](https://doi.org/10.1051/0004-6361/202554098)