

Directed searches for gravitational waves from ultralight vector boson clouds around merger remnant and galactic black holes during the first part of the fourth LIGO–Virgo–KAGRA observing run

The LIGO Scientific collaboration, the Virgo collaboration, the KAGRA collaboration
(Dated: September 10, 2025)

We present the first directed searches for long-transient and continuous gravitational waves from ultralight vector boson clouds around known black holes (BHs). We use LIGO data from the first part of the fourth LIGO–Virgo–KAGRA observing run. The searches target two distinct types of BHs and use two new semicoherent methods: hidden Markov model (HMM) tracking for the remnant BHs of the mergers GW230814_230901 and GW231123_135430 (referred to as GW230814 and GW231123 in this study), and a dedicated method using the Band Sampled Data (BSD) framework for the galactic BH in the Cygnus X-1 binary system. Without finding evidence of a signal from vector bosons in the data, we estimate the mass range that can be constrained. For the HMM searches targeting the remnants from GW231123 and GW230814, we disfavor vector boson masses in the ranges $[0.94, 1.08]$ and $[2.75, 3.28] \times 10^{-13}$ eV, respectively, at 30% confidence, assuming a 1% false alarm probability. Although these searches are only marginally sensitive to signals from merger remnants at relatively large distances, future observations are expected to yield more stringent constraints with high confidence. For the BSD search targeting the BH in Cygnus X-1, we exclude vector boson masses in the range $[0.85, 1.59] \times 10^{-13}$ eV at 95% confidence, assuming an initial BH spin larger than 0.5.

I. INTRODUCTION

Ultralight bosons are a class of theoretical particles proposed in certain extensions of the Standard Model, and their discovery could address numerous unresolved questions in particle physics and cosmology (e.g., the nature of dark matter [1–4] or the strong charge-parity-problem [5–7]). Theories predict different subclasses of ultralight bosons based on spin, including scalar (spin-0) [5–9], vector (spin-1) [10–16], and tensor (spin-2) [17–22] fields. Assuming only gravitational coupling, the superradiance mechanism provides a means by which ultralight bosons may form bound states around rotating black holes (BHs) and grow into macroscopic clouds, producing quasi-monochromatic, long-duration gravitational wave (GW) signals that may be observable by ground-based detectors [23–40].

Ultralight vector bosons with mass m_V can extract rotational energy from a BH through the superradiance mechanism if the following condition is satisfied:

$$0 < \omega < m\Omega_H, \quad (1)$$

where $\omega \approx m_V c^2/\hbar$ is the angular frequency of the boson field, m is the azimuthal quantum number, and Ω_H is the angular frequency of the BH’s outer horizon. Bound states can therefore grow exponentially in time with an instability growth rate that is maximized when the boson’s Compton wavelength λ is comparable to the BH’s horizon radius r_g . For a superradiantly unstable state, the instability extracts energy and angular momentum fluxes from the BH at the horizon into the bosonic field, with the change in mass directly related to the change in angular momentum. If we assume the primary interaction of the ultralight bosons is gravitational, this results in compound field amplification and the growth of a macro-

scopic cloud that may extract as much as $\sim 10\%$ of the BH’s initial mass [23, 37, 41–48].

This exchange of energy and angular momentum continues until Eq. (1) is saturated ($\omega = m\Omega_H$), or in other words, until the BH has been spun down so much that superradiance can no longer occur. At this point, the cloud begins to deplete through GW emission, which occurs at roughly twice the boson cloud oscillation frequency, with a small positive frequency drift as the magnitude of the bosons’ binding energy to the BH decreases [30, 37–39, 49–52]. For stellar and intermediate-mass BHs and bosons with masses within the range $\sim 10^{-14}$ – 10^{-11} eV [26, 29], this emission frequency falls in the sensitive band of ground-based GW detectors. Thus, current and future ground-based GW detector networks [53–59] offer a unique way to search for these ultralight bosons. If no detection is made, we are able to place constraints on the existence of ultralight bosons within the above mass range.

Many observational studies have already been designed and implemented to search for ultralight scalar bosons, and, in the absence of a detection, constrain their existence. Although a handful of studies have derived constraints on vector bosons as well, the vector signal morphology has only recently been accurately modeled via numerical calculations in both the relativistic and non-relativistic regimes,¹ enabling more accurate predictions of the signal morphology and more robust constraints on the boson mass [39]. Constraints are placed on the existence of ultralight scalar [26, 29, 32, 60] and vector [30, 32, 61] particles using BH spin measurements. In

¹ In the relativistic regime, the BH and the superradiant cloud are approximately the same size, whereas in the non-relativistic regime, the cloud is much larger than the BH.

some cases, however, there are systematic uncertainties related to the BH parameters that limit these studies. Searches for continuous gravitational waves (CWs) that target either the galactic center or the entire sky are used to constrain scalars and, based on certain astrophysical assumptions, they disfavor the mass range of approximately $\sim 10^{-13}$ – 10^{-12} eV [62–66]. Using the null results from searches for a stochastic GW background generated by a population of BHs with scalar [67, 68] and vector [69] clouds, boson masses $\sim 10^{-13}$ eV have been disfavored for both scalars and vectors, based on certain assumptions regarding the BH population and spin distributions. A directed search targeting the black hole in the X-ray binary Cygnus X-1 excludes scalars in the mass range of $\sim [0.6, 1] \times 10^{-12}$ eV, depending on the BH’s estimated age [70, 71]. In addition to indirect searches, ground-based interferometers are used as particle detectors for the direct detection of ultralight scalars in Refs. [72–74] and vectors in Refs. [75, 76]. Looking to the future, the existence of a scalar cloud could impact the inspiral of binary BHs, and this may soon be detectable for certain boson masses and field strengths [33, 77, 78].

In this study, we present the results from the first directed searches for GW signals from ultralight vector boson clouds (VBCs) around selected known BHs, performed using data from the first part of the fourth observing run (O4a) of the LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration (LVK). We use two semicoherent methods to target vector boson signals produced by two different categories of astrophysical sources: i) using methods outlined in Refs. [79, 80] based on a hidden Markov model (HMM), we undertake the first ever directed searches for GWs from VBCs around compact binary merger remnant BHs [81–83], and ii) using a new method based on the so-called Band Sampled Data (BSD) framework [84], we search for GWs from a VBC around the BH in the known binary system Cygnus X-1. Each method is tailored to track signals from a different type of source; while merger remnants are isolated and are expected to emit signals with rapidly evolving frequencies, Cygnus X-1 is a binary system with a BH that would emit a nearly monochromatic signal and has uncertain orbital parameters that must be accounted for. For a more in-depth discussion of these sources and their differences, see Sec. II C.

One benefit of targeting known BHs—whether they are remnants from previous GW observations or identified via X-ray emissions—is that more robust constraints can be derived. In contrast to the broader searches mentioned above, which rely on assumptions about the underlying BH population, directed searches benefit from more complete information about the target BHs. In particular, uncertainties in the BH age and spin—which can strongly influence the resulting constraints—are reduced or eliminated when the source is well-characterized. For young merger remnants, very few hypotheses are needed on the history and evolution of the BHs since their formation. On the other hand, targeting a galactic BH like Cygnus

X-1 is beneficial due to its proximity [85]. In addition, because the system is much older [86], any signal observable during O4a must correspond to an emission process that occurs over much longer timescales and produces a more slowly evolving signal (i.e., like a traditional CW), and hence the search can be run across the full O4a dataset.²

The organization of the paper is as follows. In Sec. II, we identify the target BHs and give an overview of the GW signal parameters relevant to the searches. In Sec. III, we outline the two search methods used in this study. We describe the search configurations and general setup for each method in Sec. IV. In Sec. V, we discuss the results from the searches and explain the candidate follow-up process. Then, in Sec. VI, we estimate the range of vector boson masses that can be disfavored given that no signal is detected. We summarize our findings and conclude in Sec. VII.

II. TARGET SOURCES

In this section, we give an overview of the BHs targeted by the two search pipelines, as well as their expected GW signal parameters, which can be predicted using a combination of analytic [30, 87–90] and numerical [32, 36–39, 91] methods. We generate the signal parameters using the waveform model **SuperRad**, which models the dynamics, oscillation frequency, and GW emission of ultralight boson clouds with high accuracy [52, 92].

A. Merger remnant black holes

For the HMM search, we target two binary merger remnant BHs from O4a [83]: the remnants of the GW230814_230901 [83] and GW231123_135430 [93] mergers (henceforth referred to as GW230814 and GW231123, respectively). The estimated median parameters for each remnant are shown in Table I, where M_i and χ_i are the mass and dimensionless spin of the BH before superradiance has occurred, D_L is the luminosity distance, ι is the inclination angle, and RA and Dec are the right ascension and declination of the BH. The remnants from these two events are chosen in particular because, for each BH, we find that the median estimated luminosity distance is less than the furthest reachable distance (see Fig. 9 in Ref. [79]) for a BH with the same median mass and spin. Note, because we use an earlier version of the parameter estimates (which were the most recent and accurate estimates available at the time of this study), the values shown in Table I vary slightly from those reported in Refs. [83] and [93]. However, these differences

² While merger remnants could produce similarly long-lived GW emission if the bosons with appropriate masses exist, the growth time of the corresponding clouds would far exceed the O4a time frame.

are minor and have no noticeable impact on the analyses or conclusions, as the searches are relatively insensitive to small deviations in the parameter estimates.³

The parameter estimates used in this study have been generated with the **NRSur7dq4** waveform model [94]. This model has been chosen because it interpolates between numerical relativity data without making additional waveform modeling assumptions, and it typically performs well for signals from higher-mass sources like GW231123 [93]. In standard astrophysical formation scenarios of binary systems, the high spin for the primary BH favored by the analysis of GW231123 [93] is in tension with the existence of a vector boson with mass that would give the loudest signal for the remnant BH [61]. Nonetheless, there is value in targeting the remnant for a direct and independent search.

B. Cygnus X-1

In the BSD-based search, we target the BH in the Cygnus X-1 system, with parameters listed in Table I taken from a recent study [85]. Cygnus X-1 is a binary system with an orbital period $P = 5.599829 \pm 0.000016$ days [95]. Several of its orbital parameters are constrained through X-ray observations and are summarized in Table II. The age of the BH is estimated to be $6.2 \pm 1.8 \times 10^6$ yrs [86].

Given the system's age, any detectable signal is expected to exhibit high stability, with only small variations in the emission frequency over time. This expectation is supported by simulations using the **SuperRad** waveform model, showing that the signal frequency drift remains within $\mathcal{O}(10^{-12})$ Hz s⁻¹ across the entire parameter space, which is well below the frequency resolution of the analysis. Consequently, we assume a monochromatic emission for the whole duration of O4a, and we neglect frequency derivatives typically included in standard Taylor expansions (e.g., [96, 97]).

In the literature, several studies report an extreme spin (≥ 0.95) for the BH in Cygnus X-1 [98–102]. These results disfavor the existence of a VBC since a significant part of the rotational energy of the BH would have been extracted through superradiance, so such a high spin could not be measured after superradiance took place. However, these measurements are impacted by systematic uncertainties and may depend on the accretion model. This is reflected by a disagreement in the literature with other studies

allowing spin values between 0.5 and 0.9 [103–108], some even compatible with a spin below 0.2 [109, 110]. In the Cygnus X-1 search, we ignore this tension by treating the measurable final spin of the BH after superradiance as a free parameter. To estimate the VBC properties, we adopt an initial spin value $\chi_i = 0.95$, as shown in Table I. In Sec. VI, we derive constraints on the possible boson mass for different assumed values of χ_i .

C. GW signal parameters

Using the **SuperRad** waveform model [52, 92] and the values for M_i and χ_i shown in Table I, we find the optimal vector mass—that is, the mass m_V^{opt} that optimizes superradiant instability for a given BH, producing the maximum GW strain amplitude attainable by the system at a reference epoch t_{ref} . For the remnants from GW230814 and GW231123, t_{ref} is fixed at the time when the VBC reaches its full size [79], corresponding to the time t_{sat} when Eq. (1) is saturated. In the case of Cygnus X-1, we consider the start of the O4a period as the reference time. The values of m_V^{opt} are shown for each source in the second column of Table III. Then, based on these values and the median BH parameters, we use **SuperRad** (for the fastest growing, initially dominant mode $m = 1$) to estimate the GW signal parameters for each BH. The parameters, also listed in Table III, are defined as follows. τ_{growth} and τ_{GW} are the VBC growth and depletion timescales. The parameters h_0^{eff} , f_0 , and \dot{f}_0 are, respectively, the peak GW strain amplitude in the detector frame,⁴ the GW emission frequency, and the frequency drift, each evaluated at the reference time ($t = t_{\text{ref}}$).

Here, we provide a depiction of how the signal frequency evolves once the VBC is saturated, independent of the specific target being searched. At t_{sat} , the VBC is oscillating around the BH with angular frequency $\omega = m\Omega_H$ [38]. Then, the initial frequency of the GW emission in the source frame depends on this angular frequency at t_{sat} : $f_0 = \omega/\pi$. The evolution of the GW frequency can be approximated as [52]

$$f_{\text{GW}}(t) = f_{\infty} - \frac{|f_{\text{shift}}|}{1 + (t - t_{\text{sat}})/\tau_{\text{GW}}}, \quad (2)$$

where f_{∞} is the asymptotic frequency at late times ($t - t_{\text{sat}} \gg \tau_{\text{GW}}$), and f_{shift} is the negative shift of f_{GW} away from f_{∞} as a result of the self-gravity of the VBC at t_{sat} [52, 79, 92], i.e., we have $f_0 = f_{\infty} - |f_{\text{shift}}|$. While Eq. (2) provides a useful illustration of how the emission frequency evolves over time, the frequencies and their time derivatives shown in Table III are computed with greater accuracy using **Superrad**, which incorporates

³ For GW230814, the posterior distributions of the BH properties from the preliminary version of the parameter estimation are consistent with those reported in Ref. [83]. For GW231123, the posterior distributions show minor differences compared to those presented in Ref. [93]. In particular, the distribution of D_L in Ref. [93] is more tightly constrained than the preliminary version used in this study; nonetheless, the two distributions remain consistent at the 90% confidence level. Since we adopt the broader D_L distribution, our results can be considered conservative.

⁴ h_0 in the detector frame is scaled by the orientation angle of the BH ι as follows: $h_0^{\text{eff}} = h_0 2^{-1/2} \{[(1 + \cos^2 \iota)/2]^2 + \cos^2 \iota\}^{1/2}$ [111].

Table I. Estimated median parameters for the remnant BHs from the GW230814 and GW231123 mergers and for the galactic BH Cygnus X-1.

| Source | M_i [M_\odot] | χ_i | D_L [Mpc] | $\cos \iota$ | RA [rad] | Dec [rad] |
|-----------------------|--------------------------------------|------------------------|---------------------------------|---------------------------|-------------------------------------|-------------------------------------|
| Remnant from GW230814 | $58.9^{+1.8}_{-1.8}$ ^a | $0.68^{+0.01}_{-0.02}$ | 301^{+171}_{-138} | $0.03^{+0.74}_{-0.67}$ | $3.21^{+2.62}_{-2.95}$ ^b | $0.04^{+1.02}_{-1.08}$ ^b |
| Remnant from GW231123 | $219.8^{+22.6}_{-46.2}$ ^a | $0.85^{+0.05}_{-0.18}$ | 2054^{+2960}_{-1280} | $0.45^{+0.45}_{-1.23}$ | $3.37^{+1.67}_{-0.59}$ ^b | $0.38^{+0.40}_{-0.63}$ ^b |
| BH in Cygnus X-1 | $21.2^{+2.2}_{-2.3}$ | 0.95 ^d | $0.00222^{+0.00018}_{-0.00017}$ | $0.887^{+0.005}_{-0.006}$ | 5.22883712 ^c | 0.61438355 ^c |

^a We use M_i and χ_i to represent the mass and dimensionless spin of the BH before superradiance occurs, i.e., the final BH mass (M_f) and spin (χ_f) used in the compact binary coalescence parameter estimation. The reason we do not use M_f and χ_f here is to avoid confusion with the final BH mass and spin after superradiance occurs.

^b See Fig. 1.

^c At reference epoch MJD 56198.

^d The spin of Cygnus X-1 is debated in the literature. In our analysis, we treat the final spin of the BH as a free parameter and assume a nominal initial spin of $\chi_i = 0.95$. The impact of the initial spin on the derived constraints is discussed in Sec. VI.

Table II. Orbital parameters for Cygnus X-1.

| Parameter | Symbol | Value | Ref. |
|---------------------------|----------|------------------------------|------|
| Orbital Period [days] | P | 5.599829(16) | [95] |
| Proj. Semi-major Axis [s] | a_p | $36.88^{+4.02}_{-3.65}$ | [85] |
| Eccentricity [-] | e | $0.0188^{+0.0028}_{-0.0026}$ | [85] |
| Arg. of periastron [deg] | ω | $306.6^{+6.6}_{-6.3}$ | [85] |

higher-order relativistic corrections [92]. Note that the signal characteristics differ between merger remnants and Cygnus X-1, as the depletion timescales τ_{GW} considered are different by orders of magnitude.

As is clear from Tables I and III, the characteristics of these target BHs are quite different. While the merger remnants are young, isolated BHs that are expected to emit short-duration, more rapidly evolving signals, Cygnus X-1 is a binary system with a much older BH that is expected to emit a long-duration, approximately monochromatic signal. Hence, we use two different search methods, each designed to exploit a different aspect of VBC signals: HMM is a flexible, less model-dependent method that is well-suited for signals with larger frequency drifts and greater uncertainties (see Sec. III A), whereas the Binary BSD-VBC search technique is designed for signals with negligible frequency drifts from sources whose orbital motion must be taken into account (see Sec. III B).

In the case of BHs from merger remnants with $D_L \sim \mathcal{O}(\text{Gpc})$, we must take into consideration the non-negligible impact of redshift on the GW emission frequency and timescale. The frequency in the detector frame scales as $f_{\text{det}} = f_{\text{src}}(1+z)^{-1}$, and correspondingly, the frequency derivative scales as $\dot{f}_{\text{det}} = \dot{f}_{\text{src}}(1+z)^{-2}$, where f_{src} (\dot{f}_{src}) is the frequency (derivative) in the source frame, and f_{det} (\dot{f}_{det}) is the frequency (derivative) in the detector frame. In a similar vein, the GW emission timescale is modified as $\tau_{\text{GW}}^{\text{det}} = \tau_{\text{GW}}^{\text{src}}(1+z)$. For an in-depth discussion of how these redshift corrections affect the searches, see Sec. IV in Ref. [79].

III. METHODS

In this paper, we use two semicoherent search methods: the HMM tracking method (Sec. III A) and the Binary BSD-VBC method (Sec. III B). Although both methods are based on previous work on traditional CW searches targeting individual spinning neutron stars with nonaxisymmetries [65, 84, 111, 112], this section describes how they have been tailored to searches for vector boson signals.

A. HMM method: Merger remnants

The HMM tracking technique is useful for detecting signals with wandering frequencies. It models the frequency evolution probabilistically as a Markov chain of transitions between discrete, unobservable (“hidden”) frequency states over a number of discrete time steps, and it connects these hidden states with observable data through an emission probability. HMM tracking has been used in various searches for continuous or long-transient GW signals [49, 70, 111–121]. This method is beneficial for two main reasons: i) it is computationally efficient, which is important for this study because vector boson signals live in a large, multi-dimensional parameter space, and ii) it is more capable of accounting for both the detector noise fluctuations and any uncertainties that may be present in the predicted signal waveform than other, more model-dependent CW semicoherent search techniques (see Ref. [122]).

Reference [49] was the first to propose using an HMM-based method for follow-up searches targeting slowly-evolving scalar boson clouds around BH merger remnants. Since vector bound states carry spin angular momentum, they can still grow via superradiance even with zero orbital angular momentum number ($\ell = 0$). As a result, vector bound states can be concentrated closer to the BH, leading to clouds that grow and deplete more rapidly and radiate at much higher power than scalar clouds [30, 36, 38–40]. This results in GW signals that are short-lived by CW standards [$\sim \mathcal{O}(\text{days-months})$] with frequency drift pa-

Table III. Estimated GW signal parameters for the remnant BHs from the GW230814 and GW231123 mergers and the BH in Cygnus X-1 using the median BH parameters and their respective m_V^{opt} values.

| Source | m_V^{opt} [10^{-13} eV] | τ_{growth} | τ_{GW} | h_0^{eff} [10^{-24}] | f_0 [Hz] | \dot{f}_0 [Hz s^{-1}] |
|-----------------------|-------------------------------------|------------------------|---------------------|-----------------------------------|------------|------------------------------------|
| Remnant from GW230814 | 3.805 | 7.8 h | 28 h | 0.341 | 170.4 | 6.2×10^{-7} |
| Remnant from GW231123 | 1.652 | 2.3 h | 2.5 h | 1.69 | 51.2 | 9.4×10^{-6} |
| BH in Cygnus X-1 | 1.040 | 560 y | 7.5×10^6 y | 1.75 | 50.3 | 1.6×10^{-19} |

rameters that are too large for the capabilities of the standard HMM method. Thus, building upon the work of Isi et al. [49], Ref. [79] implemented a modified HMM search pipeline capable of tracking signals from VBCs that occur over shorter timescales. The pipeline is an efficient semicoherent search method combining the HMM tracking scheme with a frequency-domain matched filter (\mathcal{F} -statistic), which quantifies the likelihood that a signal, parameterized by its frequency and associated time derivatives, is present in the data [123]. Previous studies have shown that this pipeline is sensitive enough to potentially detect signals from VBCs using data from current-generation detectors [79].

In the following two subsections, we summarize the HMM search pipeline described in Ref. [79], as well as the process of choosing search configurations for a given system, explained in Ref. [80]. Section III A 1 is meant to give a high-level overview of the HMM algorithm; for a more detailed explanation, see Refs. [111, 112].

1. Search pipeline

The search pipeline combines the HMM tracking technique with an \mathcal{F} -statistic, computed over discrete time segments. The \mathcal{F} -statistic takes as its input short Fourier transforms (SFTs) of the relevant time series data collected by each GW detector [124]. Note that for the simulated data used in various places throughout the study, we create these SFTs by injecting synthetic signal waveforms generated by the waveform model **SuperRad** into Gaussian noise using the `simulateCW` Python module in the LALPulsar library of LALSuite [125, 126].

For this pipeline, we compute the \mathcal{F} -statistic coherently over time segments of length T_{coh} . These coherent segments are then incoherently combined using an HMM, which is solved via the Viterbi algorithm [127]. The Viterbi algorithm efficiently finds the most probable path of the signal evolution (also known as the Viterbi path) in the frequency-time plane with N_Q frequency bins of width δf and N_T time segments of length T_{coh} (see detailed formulation and descriptions in, e.g., Refs. [111, 112]). The value T_{coh} (and correspondingly δf) are bounded by the predicted signal's maximum frequency derivative \dot{f}_{max} such that i) the signal may be thought of as monochromatic across a single time segment and ii) the signal can increase at most one frequency bin between each time segment. The frequency bin width δf is fixed to

$1/(2T_{\text{coh}})$. Given that the maximum spin-up of the signal across the full search duration $T_{\text{obs}} = T_{\text{coh}}N_T$ must satisfy $\dot{f}_{\text{max}}T_{\text{coh}} \leq \delta f$, we have the following upper bound on T_{coh} :

$$T_{\text{coh}} \leq (2\dot{f}_{\text{max}})^{-1/2}. \quad (3)$$

Because increasing T_{coh} generally improves the search sensitivity [111], in these searches we fix T_{coh} to its largest possible value to maximize sensitivity for a given source model configuration (see Section III A 2). The frequency drift is at its maximum when the VBC reaches saturation at a time t_{sat} . Thus, for a given system, we compute \dot{f}_{max} using **SuperRad** and then set $T_{\text{coh}} = (2\dot{f}_{\text{max}})^{-1/2}$. This ensures that the signal will not evolve outside HMM's tracking capabilities.

The Viterbi path returned by the search pipeline has an associated detection statistic that quantifies its significance. There are different ways to define the detection statistic (see, for example, the Viterbi score in, e.g., Ref. [111], or the log likelihood of the optimal path, \mathcal{L} , in, e.g., Ref. [128]). For the searches in this paper, we use \mathcal{L} divided by the number of coherent segments N_T [79], written as

$$\bar{\mathcal{L}} \equiv \mathcal{L}/N_T. \quad (4)$$

The main reason for this choice is because $\bar{\mathcal{L}}$ is more reliable for shorter duration CW searches, as compared to, e.g., the Viterbi score, which only remains reliable for very long duration searches where $N_Q \gg N_T$.

2. Configurations

Given a target BH, we start by finding m_V^{opt} . We then define a range of vector masses the pipeline *may* be sensitive to, $m_V \in [0.6, 1.1]m_V^{\text{opt}}$,⁵ given the BH we consider, and we choose some number of evenly spaced m_V values from within this range [80]. We must then choose a set of search configurations to sufficiently cover this parameter space. As demonstrated in Ref. [80], because we use the flexible HMM search technique, even a single

⁵ As shown in Ref. [79], these searches are in fact more sensitive to some sub-optimal vector masses than the optimal mass because, for a given BH, sub-optimal masses produce longer-lived GW emission, allowing us to extend the T_{coh} used in the search.

set of search configuration parameters can recover signals generated by systems with a range of parameters. Still, we require more than one configuration to provide adequate coverage of the full parameter space. The process of choosing these configurations is explained in detail in Appendix A of Ref. [80] and summarized here.

We first impose several limits on the allowed search configurations to ensure computational feasibility: We require the SFT length to lie within $15 \text{ sec} \leq T_{\text{SFT}} \leq 30 \text{ min}$ (they need not be fixed to the standard $T_{\text{SFT}} = 30 \text{ min}$ used in most CW searches) and the coherent length within $1 \text{ min} \leq T_{\text{coh}} \leq 10 \text{ day}$. The ratio $T_{\text{coh}}/T_{\text{SFT}}$ must be an integer value of at least four; that is, each detector must contribute a minimum of four SFTs per T_{coh} segment. Finally, we set $T_{\text{obs}} = \tau_{\text{GW}}$ but require that it must not exceed 180 days to minimize the computational cost of the search [79]. The search computing cost for a given system scales with the duration of the signal, ranging from approximately 10 min to 1 hr on a single-core computer. For each remnant BH for the full parameter space, the search takes roughly $\mathcal{O}(10^2)$ core-hours. If no signal is detected, deriving the final sensitivity across the full parameter space via simulations (using the same configurations) typically takes $\mathcal{O}(10^4)$ core-hours per target.

Following the guidelines outlined in Ref. [80], we randomly draw 200 posterior samples from the remnant BH's multidimensional posterior distribution. For each sample BH and each value of m_V we consider in the search, we find the optimal search configuration $\{t_{\text{start}}, T_{\text{coh}}, T_{\text{obs}}\}$, where t_{start} is the GPS start time of the search (corresponding to when the VBC has reached saturation). We then independently draw 11 values from each of these three distributions at the following percentiles: 2, 10, 20, 30, 40, 50, 60, 70, 80, 90, and 98. This forms 11 search configurations with which we will run the search (see Tables IV and V for the sets of configurations used in this work). For most systems drawn from the remnant's posterior distribution, more than one of these configurations should be able to recover the signal, making this spacing a conservative choice.

B. Binary BSD-VBC method: Cygnus X-1

The Binary BSD-VBC pipeline developed in this work builds on the BSD framework [84], which provides a compact and flexible format to manipulate and analyze calibrated strain data. Thanks to its modularity, the BSD framework has been applied to various CW searches—both fully coherent [129, 130] and semicoherent [62, 65, 75, 96, 114, 131, 132]—including the all-sky search for scalar boson clouds using data from the LVK's third observing run [62].

The Binary BSD-VBC method adapts and extends these pipelines to target CW signals emitted by VBCs around BHs in known binary systems. The main case study is Cygnus X-1, whose orbital parameters are well constrained by X-ray observations [85, 86, 95]. Details of

the specific implementation for this source are given in Sec. IV B.

1. Search pipeline

In this search, data are analyzed in sub-bands of 1 Hz overlapped by 0.5 Hz using the BSD format. The files contain a complex time series downsampled to 1 Hz and covering the full O4a.

The search pipeline follows a two-step approach based on standard BSD tools. The first step aims to increase the signal coherence by coherently removing the Doppler modulation caused by the combined motion of the source and the detector. This correction is performed, for a given set of parameters Λ , using a *heterodyne* method [84].

The second step consists of a standard semicoherent search for the demodulated signal. A collection of significant peaks in the time-frequency plane, known as the *peakmap*, is generated using the method detailed in Ref. [133]. Under proper demodulation, the signal appears in the peakmap as a line of constant frequency, with all its power contained in the same frequency bin of width $\delta f_0 = 1/T_{\text{coh}}$. To identify the signal, the peakmap is projected onto the frequency axis to produce a histogram of peak counts per frequency bin.

The significance of the number of peaks in each frequency bin is evaluated by the robust Critical Ratio (CR) statistic [97], defined in Appendix A 3. Outliers are then selected by uniformly dividing the 1 Hz band into 50 sub-bands of 0.02 Hz and choosing the frequency bin with the highest CR statistic within each sub-band.

2. Configurations

We adopt a signal model, described in Appendix A 1, where the frequency modulation is entirely defined by a set of modulation parameters $\Lambda = \{f_0, \text{RA}, \text{Dec}, a_p, t_{\text{asc}}, \Omega, e, \omega\}$ where f_0 is the emission frequency, $a_p = a \sin(\iota)/c$ with a the semi-major axis, ι the inclination angle, and c the speed of light, t_{asc} is the time of ascending node, Ω is the orbital angular frequency related to the orbital period P as $\Omega = 2\pi/P$, e is the orbital eccentricity, and ω is the argument of periape. The sampling of parameters Λ must ensure adequate coverage of the parameter space to avoid significant loss in the detection sensitivity. At the same time, the computational cost of the analysis scales with the number of templates, which must therefore be kept to a minimum. The number of templates needed to cover a parameter space is discussed in Sec. A 4.

The dimensionality of the parameter space to be covered can be reduced based on the following considerations. We substitute the unknown emission frequency with the central frequency of the 1 Hz analysis band. The search targets a system whose sky position is known with high

precision; we therefore fix RA and Dec to their electromagnetic estimates. For the orbital parameters, we further assume that the angular frequency Ω , eccentricity e , and argument of periastron ω are well-constrained from electromagnetic observations, allowing the use of single central values of Ω , e , and ω for the heterodyne correction.

With these assumptions, the parameter space for a given target BH is reduced to a two-dimensional plane (a_p, t_{asc}) . The time of ascending node t_{asc} is bounded by the orbital period as $t_{\text{asc}} \in [-P/2, P/2]$, and we further impose constraints on a_p consistent with electromagnetic observations of the source. The validity of these assumptions must be checked for each potential target of the method. The case of Cygnus X-1 is discussed in Sec. IV B and has been verified through simulated signal injections into simulated data that mimics the detectors' noise levels.

To cover the remaining parameter space, we use a restricted version of the binary search metric described in Ref. [134]. From this metric, and allowing for a maximal loss in signal-to-noise ratio of 10%, the resolutions in the orbital parameters are given by

$$\delta a_p = \frac{\sqrt{0.6}}{\pi \Omega T_{\text{coh}} f_0} \quad \delta t_{\text{asc}} = \frac{\sqrt{0.6}}{\pi \Omega^2 a_p T_{\text{coh}} f_0}. \quad (5)$$

These resolutions are used to construct a square lattice \mathbb{Z}_2 , as described in Refs. [134, 135]. Details on the grid construction can be found in Appendix A 4.

Typically, an analysis satisfying these limitations would take $< \mathcal{O}(10^5)$ core-hours to cover the reduced parameter space, such as the one of Cygnus X-1. In the absence of detection, constraints can be derived very efficiently using the method outlined in Sec. VIB 1, and the computation time needed to derive these constraints is negligible compared to the analysis time.

3. Candidate selection and coincidences

For every 1 Hz band, the search is repeated for all templates Λ , each of them producing 50 triggers. On these triggers, we select a subset of the most significant outliers. Specifically, for each frequency bin of width $\delta f_0 = 1/T_{\text{coh}}$, we select the two outliers with the highest CR values. In this way, and for the coherence time $T_{\text{coh}} = 1000$ s considered in this search, a maximum of 2000 outliers are selected for each 1 Hz band. This step reduces the number of outliers to a manageable level. Since the upper limits presented in Sec. VIB 1 are evaluated every 1 Hz based on the loudest outlier, this selection does not impact the constraints placed by this search.

Candidates are then filtered, keeping only those with a CR above a threshold. Similar to Ref. [114], the threshold is chosen for each 1 Hz band as the mean CR plus two standard deviations of the CR distribution of the outliers. This distribution is built by excluding candidates associated with known instrumental lines (see Sec. VB). We then identify pairs of outliers coincident between the

Hanford and Livingston detectors. Outliers are considered coincidental if they are on the same or adjacent frequency bins. Furthermore, coincident outliers must have compatible orbital parameters, i.e.,

$$d_{\text{metric}} = \sqrt{\left(\frac{\Delta a_p}{\delta a_p}\right)^2 + \left(\frac{\Delta t_{\text{asc}}}{\delta t_{\text{asc}}}\right)^2} < 3, \quad (6)$$

where Δa_p and Δt_{asc} denote the differences between the parameters of the candidates in each detector.

Coincident pairs of candidates, referenced hereafter as stage-2 candidates, are further analyzed in the follow-up procedure presented in Sec. VB.

IV. SEARCH SETUP

In this section, we provide details on the parameters and configurations used to run the HMM and BSD searches. We use data taken by the two Advanced LIGO detectors, Hanford and Livingston [53], during O4a, which ran from 15:00 UTC on May 24, 2023 to 16:00 UTC on January 16, 2024 [56, 136–138]. The HMM searches analyze data spanning two different time segments (corresponding to the weeks following both the GW230814 and GW231123 binary merger events), whereas the BSD search uses the full O4a data, for which the duty factors are 67% and 69% for Hanford and Livingston, respectively. All the data used in this paper, acquired when the detectors were in science observing mode, are online calibrated (i.e., low-latency C00 frames), and analysis-ready (channel names: H1:GDS-CALIB_STRAIN_CLEAN_AR and L1:GDS-CALIB_STRAIN_CLEAN_AR) [137, 139–145]. Calibration uncertainties in the strain data can affect boson parameter estimates (if a signal is detected) and influence constraints or sensitivity estimates (if not). In O4a, the $1\text{-}\sigma$ frequency- and time-dependent uncertainties are $\lesssim 10\%$ in magnitude and $\lesssim 10$ deg in phase and differ between LIGO sites [139]. However, their overall impact is subdominant to noise fluctuations. We therefore do not explicitly include calibration uncertainties in our analysis. The HMM searches take as input SFTs generated after applying a glitch gating procedure [146]. The BSD are generated from the Short FFT Database (SFDB) [133], and cleaned using the double-gating procedure described in Refs. [84, 147].

A. HMM searches: Merger remnants

1. GW230814

Recalling Sec. IIC and Table III, for the remnant BH from GW230814, we find $m_V^{\text{opt}} = 3.805 \times 10^{-13}$ eV. Thus, the interesting vector mass range $[0.6, 1.1] m_V^{\text{opt}}$ (see Sec. III A 2) becomes $[2.283, 4.185] \times 10^{-13}$ eV for this search. Each boson mass within this range, if it exists, would emit GWs at a different frequency, so we have

a corresponding frequency range of [97, 196] Hz across which we run the search, divided into 1 Hz sub-bands.




Figure 1. Joint RA and Dec posterior distribution for the merger remnants from GW230814 (top) and GW231123 (bottom). The black contours show the 90%, 50%, and 10% confidence intervals. The red crosses indicate the sky positions targeted in the search for each remnant.

When choosing which sky position(s) to target, it may seem reasonable to simply target the median RA and Dec listed in the table. However, because the estimated posterior distribution of the RA for this system is bimodal, the median is not representative of the data (see top panel of Fig. 1).⁶ Instead, we choose the optimal sky positions empirically by drawing a joint random sample of 100 RA and Dec values from the BH posterior distribution, injecting them into Gaussian noise, and attempting to

recover them using a grid of sky positions. The spacing of this grid reflects the size of the effective point spread function (EPSF) shown in Fig. 2, in which a signal has been injected at RA, Dec = [3.854, 0.370] rad and then recovered using a grid of sky positions offset from this sky position. Overall, the sky positions that recover the largest number of randomly sampled injections are RA, Dec = [1.088, -0.387], [1.388, -0.387], [3.254, 0.370], and [3.554, 0.370] rad. We target all four sky positions in the search to obtain better coverage over the whole sky.⁷ They are marked with red crosses in the top panel of Fig. 1. The search is not particularly sensitive to sky localization for short-duration signals with $\tau_{\text{GW}} \lesssim 1$ day, although the low sky resolution does lead to degraded sensitivity [79]. For longer signals with $\tau_{\text{GW}} \gtrsim 1$ day, however, a mismatch in sky position can still cause a marginal signal to be missed. This uncertainty due to poorly-constrained sky localization is incorporated into the sensitivity estimates we present in Sec. VI A.




Figure 2. Colored contour of $\bar{\mathcal{L}}/\bar{\mathcal{L}}_{\text{th}}$ as a function of RA and Dec for a synthetic signal injected at [3.854, 0.370] rad, shown with a white cross marker. The signal was generated using the median BH remnant parameters from the GW230814 merger shown in Table I and using $m_V = 2.283 \times 10^{-13}$ eV. The bright EPSF enclosed within the white contour signifies the region of the sky with $\bar{\mathcal{L}} > \bar{\mathcal{L}}_{\text{th}}$ where the signal has been recovered.

Using the method outlined in Appendix A of Ref. [80] for choosing search configurations, we use 11 different configurations in the search, shown in Table IV. In particular, we choose a set of T_{coh} values to cover the full range of f_{max} values that may occur for the system: $\sim \mathcal{O}[10^{-9}, 10^{-6}]$ Hz s⁻¹. Across all 11 configurations,

⁶ GW230814 was observed only by the Livingston detector, so the remnant's poorly constrained sky location aligns with the detector's antenna pattern.

⁷ It is particularly valuable to target multiple sky positions (when the true sky position is not well constrained) for the configurations that use longer T_{coh} values. This is because searches with longer coherent segments are more sensitive to Doppler modulation effects and thus the assumed sky position, so there is a larger chance of missing a GW signal if an incorrect sky position is used.

we run the search on detector data spanning the GPS times [1376111180, 1385937626] s. Gaps are present in the available data from Hanford and Livingston during this time period, resulting from both scheduled maintenance and unexpected lock loss due to various environmental disturbances. In particular, although the duty factor is $\sim 70\%$ for each detector during the analysis time of GW230814, many of the data gaps lie within roughly two days post-merger, when the signal is expected to be strongest. While the HMM search pipeline is designed to accommodate data gaps, the less data that is available, the less sensitive the search becomes. This decreased sensitivity is reflected in the eventual estimated vector mass range that is disfavored by targeting this remnant (see Sec. VIA).

2. GW231123

As in the previous section, we refer to Table III for the optimal boson mass corresponding to the remnant BH from GW231123: $m_V^{\text{opt}} = 1.652 \times 10^{-13}$ eV. Thus, for the search we consider the range $m_V = [0.901, 1.502] \times 10^{-13}$ eV and the corresponding frequency band [21, 73] Hz, split into 1 Hz sub-bands.

Once again, the RA posterior distribution shows some bimodality, but because the sky position is significantly more well-constrained than GW230814 [i.e., the 90% confidence interval is constrained to $\mathcal{O}(10^3)$ deg²], for the search we simply target the two local maxima marked by red crosses in the bottom panel of Fig. 1, located at RA, Dec = [3.329, 0.372] and [5.153, 0.318] rad. These two local maxima are tested using the same method described above for GW230814, and we find that they provide sufficient coverage across the 90% credible region of the sky position.

We start by identifying 11 potential search configurations at the same percentiles used for GW230814. However, two configurations cannot be used in the search, as their GPS start times (1602835969 and 2785340449 for the 90th and 98th percentiles, respectively) do not fall within the fourth observing run (O4).⁸ Thus, we limit the number of configurations used in the search to only the first 9, shown in Table V. Again, we choose the set of T_{coh} values to cover the potential \dot{f}_{max} range $\sim \mathcal{O}[10^{-15}, 10^{-5}]$ Hz s⁻¹ for this system. Gaps are again present in the data available from both detectors, limiting the search sensitivity; while the duty factor for Hanford is 68% across the relevant time frame, for Livingston it is only 51%.

Table IV. Search configuration parameters and the percentiles at which they are drawn for the BH remnant from the GW230814 merger. These percentiles are chosen so that the configurations used in the search adequately cover the remnant’s full posterior distribution.

| Percentile | T_{SFT} [m] | T_{coh} [m] | T_{obs} [m] | GPS start time [s] |
|------------|----------------------|----------------------|----------------------|--------------------|
| 2 | 2.85 | 11.4 | 991.8 | 1376111180 |
| 10 | 3.4 | 13.6 | 1278.4 | 1376115080 |
| 20 | 3.9 | 15.6 | 1653.6 | 1376119160 |
| 30 | 4.55 | 18.2 | 2311.4 | 1376124380 |
| 40 | 5.75 | 23.0 | 3519.0 | 1376134280 |
| 50 | 7.45 | 29.8 | 5632.2 | 1376148320 |
| 60 | 10.45 | 41.8 | 9823.0 | 1376172680 |
| 70 | 15.05 | 60.2 | 17759.0 | 1376211140 |
| 80 | 24.4 | 97.6 | 36990.4 | 1376285360 |
| 90 | 26.85 | 161.1 | 79583.4 | 1376412320 |
| 98 | 27.65 | 248.85 | 155780.1 | 1376590820 |

Table V. Search configuration parameters and the percentiles at which they are drawn for the BH remnant from the GW231123 merger. These percentiles are chosen so that the configurations used in the search adequately cover the remnant’s full posterior distribution.

| Percentile | T_{SFT} [m] | T_{coh} [m] | T_{obs} [m] | GPS start time [s] |
|------------|----------------------|----------------------|----------------------|--------------------|
| 2 | 0.6 | 2.4 | 84.0 | 1384788409 |
| 10 | 1.0 | 4.0 | 172.0 | 1384792009 |
| 20 | 1.45 | 5.8 | 301.6 | 1384796449 |
| 30 | 2.25 | 9.0 | 630.0 | 1384804309 |
| 40 | 3.55 | 14.2 | 1235.4 | 1384816729 |
| 50 | 6.0 | 24.0 | 2424.0 | 1384837489 |
| 60 | 9.4 | 37.6 | 4587.2 | 1384866949 |
| 70 | 14.35 | 57.4 | 9011.8 | 1384911769 |
| 80 | 26.2 | 157.2 | 42601.2 | 1385145169 |

B. BSD search: Cygnus X-1

Using the median values for the parameters of Cygnus X-1 reported in Table I, we find the boson mass $m_V^{\text{opt}} = 1.040 \times 10^{-13}$ eV produces the signal with the maximum strain amplitude at the start of O4a. Taking into account the uncertainties in the BH parameters, we estimate the signal amplitudes corresponding to boson masses around m_V^{opt} using the SuperRad model. Then, comparing these values to the minimum detectable strain given in Eq. (67) of Ref. [97], we select the frequency band $f_0 \in [24.5, 125.5]$ Hz as the range of frequencies where a signal could be detected. Therefore, the search spans 202 distinct 1 Hz bands, overlapped by 50%.

We fix the coherence time to $T_{\text{coh}} = 1000$ s for all frequency bands investigated. This value is motivated by multiple factors. Demodulating the data using the central frequency of the band rather than the real—unknown—

⁸ These times correspond to a non-optimal region of the parameter space where the VBC grows and dissipates very slowly and thus it will be $\mathcal{O}(\text{years})$ before the GW emission reaches its peak.

frequency of the signal leads to a residual modulation of the signal. Our choice of T_{coh} ensures that, for any signal present in the band, the residual modulation Δf is smaller than the size of a frequency bin, i.e. $2\Delta f < 1/T_{\text{coh}}$. Similarly, this coherence time ensures that the uncertainties in the eccentricity parameters of Cygnus X-1's orbit, reported in Table II, are fully covered by a template using the central values of the parameters. At the same time, a coherence time $T_{\text{coh}} = 1000$ s keeps the analysis time to a realistic level with $\mathcal{O}(60\,000)$ core-hours being needed to analyze the parameter space of Cygnus X-1 in both LIGO detectors.

Following the grid construction method presented in Sec. III B 2, for each band we define a set of parameter points $\{\mathbf{A}_i\}_{i=1}^N$ covering the uncertainties in the orbital parameters of Cygnus X-1 reported in Table II. The time of superior conjunction T_0 of Cygnus X-1 has been estimated in Ref. [95], but the extrapolation of this value to estimate the time of ascending node t_{asc} at the O4a period leads to significant uncertainties compared to our search resolution. We choose an agnostic approach where the range of all possible values $t_{\text{asc}} \in [-P/2, P/2]$ has been covered by the search. The number of templates per band ranges from $N = 1960$ at 24.5 Hz to $N = 43\,416$ at 125.5 Hz, in agreement with Eq. (A12). Combining the number of templates used in all the bands, the search uses 3 687 225 templates.

The search is performed with bands overlapped by 0.5 Hz to avoid cropping signals close to the band edges during BSD creation. Each band is searched independently, and each produces its own set of outliers, among which outliers and stage-2 candidates are selected following the discussion in Sec. III B 3.

V. CANDIDATE FOLLOW-UP

In this section, we outline the follow-up procedures used to eliminate any search candidates whose origins are not astrophysical. See Tables VI and VII for the numbers of candidates that remain from the HMM and BSD searches after each veto procedure described in the following sections.

A. HMM searches: Merger remnants

For GW230814, we run the search over 99 individual 1 Hz bands, 11 search configurations, and 4 sky positions. This is 4356 iterations in total. Similarly, for GW231123, we have 52 1 Hz bands, 9 search configurations, and 2 sky positions, yielding 936 search iterations. For each 1 Hz band, we require the detection statistic $\tilde{\mathcal{L}}$ to exceed a threshold $\tilde{\mathcal{L}}_{\text{th}}$ corresponding to a 1% false alarm proba-

bility (P_{fa}).⁹ The value of $\tilde{\mathcal{L}}_{\text{th}}$, which varies for each T_{coh} used in the searches, is obtained empirically as follows: For a given value of T_{coh} , we run 300 searches in pure Gaussian noise (with Amplitude Spectral Density (ASD) $= 4 \times 10^{-24}$ Hz $^{-1/2}$) at a randomly chosen 1 Hz frequency band, and we define $\tilde{\mathcal{L}}_{\text{th}}$ as the value of $\tilde{\mathcal{L}}$ that lies at the 99th percentile of these results.

After implementing this initial threshold cut, we have 421 and 285 signal candidates across all configurations for GW230814 and GW231123, respectively. However, many candidates actually appear to be the same candidates picked up by different configurations (e.g., for GW230814, a candidate at ~ 102.1 Hz is identified for nearly every value of T_{coh} and at every sky position). We expect a large fraction of these first-pass candidates to be simply the result of noise artifacts (e.g., power line harmonics, thermally excited mirror suspension violin modes [137, 148]) and non-Gaussianities in the interferometric data. We use the known-line [149] and single-interferometer veto techniques described in Appendix B 1 (which are commonly used in CW searches) to help distinguish candidates of this nature from a true astrophysical signal [128]. We manually inspect any candidates that survive these initial vetoes by assessing their consistency with the signal model, scrutinizing the spectrograms, identifying clear characteristics of noise (such as a candidate occurring in only one detector due to short-period artifacts and failing to meet the stringent criteria of the single-interferometer veto), etc. For a detailed description, see Appendix B 1. Table VI shows the candidates that remain for each merger remnant after each veto. After manual inspection, no signal candidates remain.

Table VI. Number of candidates remaining from the search targeting the remnants from the GW230814 and GW231123 mergers after each veto has been applied.

| Stage | Remaining candidates | |
|----------------------------|----------------------|----------|
| | GW230814 | GW231123 |
| Initial candidates | 421 | 285 |
| Known-line veto | 44 | 209 |
| Single-interferometer veto | 19 | 101 |
| Manual inspection | 0 | 0 |

B. BSD search: Cygnus X-1

From the BSD-based search, we identify 27 stage-2 candidates that pass the selection process described in Sec. III B 3. We follow up these candidates using the vetoing procedure described in Appendix B 2. The list of

⁹ The 1% false alarm probability corresponds to each configuration in each 1 Hz band.

candidates is first filtered by removing outliers associated with known instrumental lines [149], reducing the list to eight candidates. During the creation of outlier pairs (see Sec. III B 3), the same outlier could be matched to multiple coincident outliers in the other detector. To remove this redundancy, we cluster all pairs sharing a common outlier, which reduces the number of independent candidates to six. We then apply standard vetoes adapted from previous BSD-based searches (e.g., [65, 114]), which are based on the consistency of the candidates with the parameters of Cygnus X-1 and with the expected behavior of the detection statistic. Three candidates pass these vetoes and are further inspected. All three are consistent with artifacts produced by non-Gaussianities in the detector. More details on the follow-up procedure are given in Appendix B 2. Table VII summarizes the number of candidates remaining after each veto.

Table VII. Number of candidates remaining from the Cygnus X-1 search after each veto has been applied.

| Stage | Remaining candidates |
|--------------------------------|----------------------|
| | Cygnus X-1 |
| Initial candidates | 27 |
| Known-line veto | 8 |
| Clustering | 6 |
| Astrophysical consistency veto | 5 |
| Statistical consistency check | 3 |
| Manual inspection | 0 |

VI. CONSTRAINTS

In this section, we estimate the vector boson mass range that can be constrained given the absence of a confident detection in the searches described above.

A. HMM searches: Merger remnants

After applying the veto procedures described in Sec. V A, all candidates from the HMM searches are eliminated. In this section, we investigate the confidence with which we disfavor the existence of a given vector boson mass range. We adopt an empirical approach in which synthetic signals are injected into simulated Gaussian noise configured to match the real data, with ASDs derived from detector data at the corresponding times and frequencies, and with data gaps reproduced to reflect those present during the analysis period. We marginalize the detection probabilities over the BH parameter uncertainties to reduce potential biases in the sensitivity estimates. We make a common assumption when interpreting the search results: the vector field interacts only gravitationally, with no additional interactions or cou-

plings to the Standard Model. See Sec. IV of Ref. [80] for details on how this assumption can be partially lifted.

In Ref. [80], a framework is developed for constraining the boson mass that marginalizes over the parameter uncertainties typical to a binary merger remnant detected gravitationally. We start by drawing a number N_{BH} of random samples from the BH posterior distribution with parameters θ_i . Then, for a given boson mass, we generate a synthetic GW signal for each sample BH and inject the signal into a number N_{noise} of random Gaussian noise simulations. We evaluate the recovery rate across all sampled systems and noise realizations as shown:

$$P_{\text{det}}(m_V) = \frac{1}{N_{\text{BH}} N_{\text{noise}}} \sum_{i=1}^{N_{\text{BH}}} N_{\text{det}}(\theta_i; m_V), \quad (7)$$

where $N_{\text{det}}(\theta_i; m_V)$ is the number of recovered ($\bar{\mathcal{L}} > \bar{\mathcal{L}}_{\text{th}}$) signals out of N_{noise} noise realizations for a given vector mass m_V and set of BH parameters θ_i . This value $P_{\text{det}}(m_V)$ can be interpreted as the confidence to which the existence of the vector boson with mass m_V can be excluded given that no signal is detected in the searches.

In Fig. 3 we show the confidence with which we disfavor the vector mass using this procedure for GW230814 (left panel) and GW231123 (right panel), where, following the guidelines in Ref. [80], we have used $N_{\text{BH}} = 200$ and $N_{\text{noise}} = 10$. We run the simulations across all configurations used in the real searches (i.e., {RA, Dec} and { T_{SFT} , T_{coh} , T_{obs} , and search start time}). The synthetic signal is considered recovered if at least one configuration returns an above-threshold detection statistic. The orange, blue, and purple lines in the figure indicate a 1%, 5%, and 10% P_{fa} threshold, respectively. P_{det} indicates the confidence with which one can disfavor a vector mass range given the null search results (i.e., if some range of data points lie above a given P_{det} value, that mass range is disfavored with P_{det} confidence). While the searches targeting the remnant from GW230814 and GW231123 are not sensitive enough to constrain the vector mass at high confidence, the following is an example of how we would estimate the boson mass range that can be constrained: We disfavor the vector mass ranges $[2.75, 3.28] \times 10^{-13}$ and $[0.94, 1.08] \times 10^{-13}$ eV for GW230814 and GW231123, respectively, with 30% confidence for $P_{\text{fa}} = 1\%$.

These results show that we are approaching the sensitivity required to place robust constraints on the vector boson mass. As the detectors continue to undergo improvements in subsequent observing runs, we anticipate a growing number of high-SNR events, which will enable increasingly sensitive searches and stronger constraints across a wide range of the mass parameter space. In addition, once the vector mass can be constrained with higher confidence, the results can be mapped to other interaction models [80].

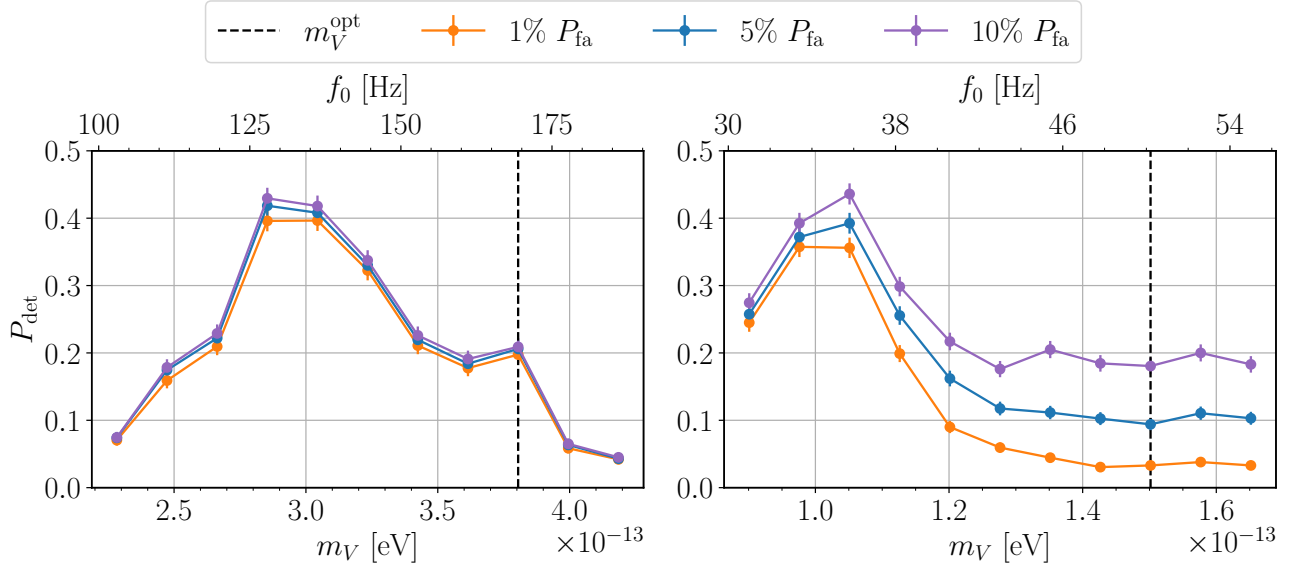


Figure 3. The detection probability P_{det} as a function of m_V (bottom axis) and the corresponding GW frequency f_0 in the detector frame (top axis) with a 1% (orange), 5% (blue), and 10% (purple) false alarm probability for the remnant BHs from the GW230814 (left) and GW231123 (right) mergers. A redshift correction has been applied to obtain f_0 in the detector frame using the median D_L value for each remnant (as reported in Table I). The vertical dashed lines mark the optimal boson masses m_V^{opt} for the remnants with median parameters also shown in Table I. The error bars represent the 1σ beta-binomial uncertainty in P_{det} [150].

B. BSD search: Cygnus X-1

Among the 27 stage-2 candidates obtained from the BSD search, none pass the follow-up procedure. In the absence of a plausible signal candidate, we set upper limits on the strain amplitude of CWs emitted by a VBC surrounding Cygnus X-1. In the following subsections, we outline the procedure used to estimate these upper limits and describe how they are subsequently translated into constraints on the mass of a hypothetical ultralight vector boson.

1. Upper limits and Sensitivity

In this section, we estimate upper limits on the strain amplitude h_0 , defined as the maximum amplitude above which the presence of a CW signal can be excluded at a given confidence level. The limits are computed using a conservative semi-analytical method previously applied in Refs. [62, 65, 96]. The method evaluates the upper limits in each 1 Hz frequency band using the analytical relation [62, 65, 96]

$$h_{\text{UL}}^{95\%} \approx \frac{\mathcal{B}}{N^{1/4}} \sqrt{\frac{S_n(f)}{T_{\text{coh}}}} \sqrt{\text{CR}_{\text{max}} + 1.645}, \quad (8)$$

where $N \sim T_{\text{obs}}/T_{\text{coh}}$ is the number of segments used to construct the peakmap, the value 1.645 is computed from Eq.(67) in Ref. [97] fixing the confidence level to

95%, $S_n^{1/2}(f)$ is the detector average ASD, and CR_{max} is the maximum CR value observed in the band. The coefficient \mathcal{B} accounts for the average detector response to the source sky position and signal polarization and depends on the peak selection threshold. For this search and the sky position of Cygnus X-1, we evaluate $\mathcal{B} \approx 3.37$ in Livingston and $\mathcal{B} \approx 3.30$ in Hanford. The computation of these factors is detailed in Appendix A 5 following the discussions in Refs. [65, 97]. The final upper limit placed in each band is taken as the less constraining limits of the two detectors.

We have verified that this semi-analytical approach yields conservative upper limits compared to those obtained with a classical frequentist approach based on simulated signal injections, while requiring fewer computational resources. This verification was performed on ten 1 Hz frequency bands, randomly selected within the search frequency range, using simulated data that reproduces the detectors' noise levels. The validation complements previous checks carried out for other observing runs of the LIGO and Virgo detectors and for similar BSD-based methods [65, 151].

In Fig. 4, we show the joint upper limits at a 95% confidence level, reporting in each frequency band the worst-case estimate between the two detectors. We also report the boson masses corresponding to the frequency axis, assuming the central astrophysical parameters of Cygnus X-1 reported in Table I.

We quantify the performance of the search using the

sensitivity depth [152], defined as

$$\mathcal{D}^{95\%} = \frac{\sqrt{S_n(f)}}{h_{\text{UL}}^{95\%}}. \quad (9)$$

This quantity has become a key figure for assessing the performance of a CW search configuration, independent of the noise level [96, 153–155]. For the BSD search, we report a sensitivity depth of $37.22 \text{ Hz}^{-1/2}$ in the band centered on 50.5 Hz corresponding to the optimal boson mass for Cygnus X-1 (see Table III). The average depth over all the analyzed bands is $35.15 \text{ Hz}^{-1/2}$ and is more or less constant in all the bands, excluding the one with significant noise disturbances.




Figure 4. Upper limit estimates at a 95% confidence level (black curve) as a function of frequency (top axis) and the corresponding boson mass assuming the central values for the mass and age of Cygnus X-1 (bottom axis). The hashed regions correspond to the predicted strain amplitude from a VBC around Cygnus X-1, assuming the central value (orange, circles), 5th percentile (blue, horizontal lines), and 95th percentile (purple, vertical lines) of the BH mass posterior distribution from [85]. We assume the BH initial spin to be $\chi_i = 0.95$ and the hashed regions account for uncertainties in the BH age and distance.

2. Constraints

The upper limits obtained in the previous section can be used to constrain the existence of ultralight vector bosons by comparing them to the expected strain amplitude of a signal emitted by a VBC in Cygnus X-1. The amplitude is estimated assuming the central value, the 5th, and the 95th percentiles of the BH mass posterior distribution reported in Ref. [85]. We also consider the uncertainties in the BH age and distance and assume the BH spin before the superradiant instability to be $\chi_i = 0.95$. The amplitude is then evaluated at different boson masses (i.e., at different frequencies) using the **SuperRad** model [52, 92].

The simulated strain amplitudes are shown in Fig. 4, along with the upper limits derived in the previous section.

By comparing the upper limits with the expected strain amplitude, we can exclude the presence of a signal emitted by a VBC in Cygnus X-1 in every frequency band in the range 41.25–80 Hz, with the exception of the narrow band 69.75–70.25 Hz. The absence of a signal in these bands excludes the existence of an ultralight vector boson with a mass in the range $[0.85, 1.65] \times 10^{-13} \text{ eV}$ (excluding the band $[1.44, 1.45] \times 10^{-13} \text{ eV}$).

These exclusion regions are obtained under the assumption that the initial spin of the BH was $\chi_i = 0.95$. This assumption can be relaxed by computing the exclusion region for different initial spin values. As the assumed initial spin decreases, the expected amplitude decreases accordingly. Nevertheless, based on **SuperRad** simulations, for any value $\chi_i > 0.5$, our upper limits can constrain the existence of a signal between 41.25 and 77.25 Hz, corresponding to boson masses within the range $[0.85, 1.59] \times 10^{-13} \text{ eV}$ (with the same excluded band as before). For spin values $\chi_i < 0.5$, the lower expected amplitude narrows down the constrained frequency range. At $\chi_i = 0.2$, we are still able to constrain frequencies in the range 50.25–77.25 Hz and boson masses $[1.03, 1.59] \times 10^{-13} \text{ eV}$. For lower initial spin values, the expected amplitude decreases rapidly, quickly leaving the existence of an ultralight vector boson unconstrained by our search. All of these exclusion intervals are computed at a 95% confidence and using the least constraining values for the BH age, distance, and mass.

VII. CONCLUSION

In this paper, we carry out the first directed searches for long-duration, quasi-monochromatic GWs from VBCs around known BHs. We analyze data from the first part of the LVK’s fourth observing run, and we use two semi-coherent CW search methods, HMM tracking and the Binary BSD-VBC pipeline. Having found no evidence of a GW signal, we estimate the range of ultralight vector boson masses that can be constrained. From the HMM search, we disfavor the vector mass ranges $[0.94, 1.08]$ and $[2.75, 3.28] \times 10^{-13} \text{ eV}$ at 30% confidence ($P_{\text{fa}} = 1\%$). While the present search sensitivity is limited because we target remnant BHs with SNR values $\lesssim 40$ —and thus the confidence level remains statistically insignificant—future searches targeting higher SNR events are expected to yield high-confidence constraints. Meanwhile, the BSD-based search excludes the mass range $[0.85, 1.59] \times 10^{-13} \text{ eV}$ at 95% confidence, assuming an initial spin value $\chi_i > 0.5$ for Cygnus X-1.

As the BSD search demonstrates, we are now able to set constraints on the existence of ultralight vector bosons by targeting known galactic BHs. Although the constraints obtained from Cygnus X-1 cover only a narrow range of masses, future improvements in detector sensitivity will enhance these constraints. The boson mass ranges that can be constrained also highly depend upon the BH

parameters. Therefore, running similar searches targeting other known galactic BHs could extend the exclusion region.

The first HMM search targeting binary merger remnants demonstrates that we are approaching the required sensitivity to place high-confidence constraints on a range of vector masses. Targeting young merger remnants allows us to set independent constraints with minimal assumptions about the BH’s history and evolution since formation. Several contributing factors will improve the sensitivity of this type of search: in particular, both increased detector sensitivity and improved search methodologies. Future observing runs and next-generation GW detectors will offer enhanced sensitivity, enabling the detection of numerous binary mergers in the high-SNR regime [57–59, 83]. These high-SNR events will yield remnant BHs with masses and spins that are more accurately and precisely measured. For HMM-based searches, improved posteriors on these parameters will directly translate into tighter constraints on the vector boson mass. In addition, future searches can incorporate modifications to the analysis pipelines to further improve search sensitivity (e.g., extending the HMM-based analysis in suitable cases to track time derivatives of the signal frequency within the \mathcal{F} -statistic, as demonstrated in Ref. [121]).

Observational studies have indirectly constrained the existence of ultralight vector bosons either from BH spin measurements [30, 32], by reinterpreting the results of GW searches [64], or from searches for a stochastic GW background [69]. A recent analysis using the GW231123 and GW190517 constituent BHs disfavors a vector mass range of $[0.11, 18] \times 10^{-13}$ eV assuming a BH age of 10^5 yrs [61]. In other searches, the interaction of ultralight vectors with ground-based GW detectors has been directly constrained [75, 76]. While each approach involves its own assumptions and limitations, the results presented in this paper provide independent constraints obtained from directed searches that, consistent with previous studies, disfavor the existence of vector bosons with masses of $\sim 1 \times 10^{-13}$ eV.

VIII. 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 - Centro Nazionale 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 Direcció General de Recerca, Innovació i Transformació Digital with funds from the Tourist Stay Tax Law ITS 2017-006, the Conselleria d’Economia, Hisenda i Innovació, the FEDER Operational Program 2021-2027 of the Balearic Islands, 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 Regional 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), the supercomputing facilities of the Université catholique de Louvain (CISM/UCL) and the Consortium des Équipements de Calcul Intensif en Fédération Wallonie Bruxelles (CÉCI), 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 Science 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 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, the University of Tokyo, the National Research Foundation (NRF), the Computing Infrastructure Project of 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 Rising Star Program and Science Vanguard Research Program, the Advanced Technology Center (ATC) of NAOJ, and the Mechanical Engineering Center of KEK.

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.

Appendix A: Details on BSD search pipeline

1. Signal phase

The phase evolution of a monochromatic signal emitted by a source in a binary system is, in the detector frame, given by

$$\Phi(t) = 2\pi f_0 t + \Delta\phi(t), \quad (\text{A1})$$

where f_0 is the signal frequency. The phase modulation due to the combined motion of the source and the detector is given by

$$\Delta\phi(t) = 2\pi f_0 \left[\frac{\mathbf{r} \cdot \hat{\mathbf{n}}}{c} - \frac{R(t + \frac{\mathbf{r} \cdot \hat{\mathbf{n}}}{c})}{c} \right], \quad (\text{A2})$$

with \mathbf{r} the position vector of the detector relative to the Solar-System Barycenter (SSB), and $\hat{\mathbf{n}}$ the unit vector pointing from the SSB to the source. The Rømer delay, R , can be expressed in a low-eccentricity orbit approximation (consistent with the nearly circular orbit of Cygnus X-1), neglecting an irrelevant constant term [134]:

$$\frac{R(t)}{c} = a_p \left[\sin(\psi(t)) + \frac{e \cos(\omega)}{2} \sin(2\psi(t)) - \frac{e \sin(\omega)}{2} \cos(2\psi(t)) \right], \quad (\text{A3})$$

where a_p is the projected semi-major axis, e is the orbital eccentricity, and ω is the argument of periape. The function ψ is the mean orbital phase measured from the time of ascending node t_{asc} , defined as

$$\psi(t) = \Omega(t - t_{\text{asc}}), \quad (\text{A4})$$

with Ω the orbital angular frequency.

2. Heterodyne correction

The procedure of the heterodyne correction of BSD is presented in Ref. [84]. It consists of multiplying the data by a complex phase factor $\exp(-i\Delta\phi(t))$ to correct Doppler-induced variations. The phase $\Delta\phi(t)$ is computed from Eq. (A2) assuming a set of parameters $\mathbf{\Lambda}$ to compensate for the delays in the signal arrival time from the motion of the source and the detector. The modulation also depends on the unknown signal frequency. Taking advantage of the BSD framework and the down-sampling of the data into 1 Hz bands, we substitute for f_0 a reference value fixed at the central frequency. By doing so, we ensure that the error in frequency is bounded by $|f_0 - f_{0,\text{ref}}| \leq 0.5$ Hz.

3. Detection statistic

The CR statistic is estimated on a peakmap by projecting the peakmap on the frequency axis and producing the distribution of number of peaks per frequency bin $n(f)$. We then use a robust estimator defined as [97]

$$\text{CR}(f) = \frac{n(f) - \bar{n}}{\sigma}, \quad (\text{A5})$$

where \bar{n} is the median of the number of peaks per frequency bin, and the dispersion parameter σ is defined by

$$\sigma = \frac{\text{median}(|n(f) - \bar{n}|)}{0.6745}. \quad (\text{A6})$$

The normalization factor σ ensures that if n follows a normal distribution, then σ is the standard deviation.

4. Template placement

For the placement of templates to cover the a_p, t_{asc} parameter space, we use a restricted version of the binary search metric described in Ref. [134]. We use the metric in the semicoherent short-segment regime, where $T_{\text{coh}} \ll P$,

$$g_{a_p a_p} = \frac{1}{6} (\pi \Omega T_{\text{coh}} f_0)^2, \quad (\text{A7})$$

$$g_{t_{\text{asc}} t_{\text{asc}}} = \frac{1}{6} (\pi \Omega^2 T_{\text{coh}} f_0 a_p)^2. \quad (\text{A8})$$

The resolution in an orbital parameter i is then given by [134]¹⁰

$$\delta_i = \sqrt{0.1[g^{-1}]^{ii}}, \quad (\text{A9})$$

where g^{-1} is the inverse metric, and the factor 0.1 corresponds to a maximal loss of signal-to-noise ratio of 10%. Explicitly, we get for the two remaining orbital parameters a_p and t_{asc} :

$$\delta a_p = \frac{\sqrt{0.6}}{\pi \Omega T_{\text{coh}} f_0}, \quad (\text{A10})$$

$$\delta t_{\text{asc}} = \frac{\sqrt{0.6}}{\pi \Omega^2 a_p T_{\text{coh}} f_0}. \quad (\text{A11})$$

Constructing a template grid with varying resolution across the parameter space can be challenging. To simplify this, we adopt a conservative strategy by fixing all the variable parameters to values that maximize grid density. Specifically, we use the maximum frequency $f_{0,\text{max}}$ within the 1 Hz band, and the maximum value of the projected semi-major axis $a_{p,\text{max}}$ within the range to be covered. Fixing the resolutions as $\delta a_p(f_{0,\text{max}})$ and $\delta t_{\text{asc}}(f_{0,\text{max}}, a_{p,\text{max}})$ over the entire parameter space ensures an overcoverage of the search parameter space. With this simplification, the template grid is constructed using a standard square lattice \mathbb{Z}_2 , as described in Refs. [134, 135].

With this placement strategy, the number of templates needed to cover the parameter space is given by [134]

$$N \approx \left\lceil \frac{\mathcal{R}_{a_p}}{\sqrt{2}\delta a_p} \right\rceil \left\lceil \frac{\mathcal{R}_{t_{\text{asc}}}}{\sqrt{2}\delta t_{\text{asc}}} \right\rceil, \quad (\text{A12})$$

where \mathcal{R}_{a_p} and $\mathcal{R}_{t_{\text{asc}}}$ are the sizes of the dimensions to be covered.

5. Upper limits formula

The average prefactor \mathcal{B} in Eq. 8 is computed by reproducing the derivation of Eq.(67) in Ref. [97] and according to the correction detailed in Ref. [156]. The general expression of \mathcal{B} is a function of time, the signal polarization angle ψ , the source sky-position (RA, Dec), and the source inclination angle ι . It can be expressed as

$$\mathcal{B} = \sqrt{\frac{\pi}{2.4308 \langle (F_+ A_+ + F_\times A_\times)^2 \rangle_t}} \left(\frac{p_0(1-p_0)}{p_1^2} \right)^{1/4} \quad (\text{A13})$$

where the factor $\frac{\pi}{2.4308}$ is taken from Eq. B18 in Ref. [97], and p_0 and p_1 are functions of the peak selection threshold

θ_{thr} [97, 156]. For $\theta_{\text{thr}} = 2.5$, we have $p_0 = 0.075$ and $p_1 = 0.096$, using the updated definition of p_1 shown in Ref. [156]. The two beam pattern functions F_+ and F_\times are defined in Ref. [123], and the polarization amplitudes are given by $A_+ = \frac{1+\cos^2\iota}{2}$ and $A_\times = \cos\iota$ [97].

For the BSD search, we evaluate the expression of \mathcal{B} for the sky-position of Cygnus X-1, and by averaging over the inclination angle error range [85] and the polarization angle $\psi \in [-\pi/4, \pi/4]$. Using properties of the beam pattern function [65], we can write

$$\begin{aligned} \langle (F_+ A_+ + F_\times A_\times)^2 \rangle_{t,\psi,\cos\iota} \Big|_{\text{Dec}} \\ = \langle F_\times^2 \rangle_{t,\psi} \Big|_{\text{Dec}} \langle A_+^2 + A_\times^2 \rangle_{\cos\iota}. \end{aligned} \quad (\text{A14})$$

We compute $\langle A_+^2 + A_\times^2 \rangle_{\cos\iota} = 1.582$, and, following the development in Ref. [65], $\langle F_\times^2 \rangle_{t,\psi} \Big|_{\text{Dec}} = 0.197$ in Livingston and 0.206 in Hanford. Injecting these values in Eq. A13 gives $\mathcal{B} \approx 3.37$ in Livingston and $\mathcal{B} \approx 3.30$ in Hanford.

Appendix B: Follow-up vetoes

As detailed in this section, we only veto candidates that we are confident are caused by noise artifacts. The safety of the vetoes used prior to the final manual inspection has been verified in previous analyses employing HMM- and BSD-based techniques, using Monte Carlo simulations in clean frequency bands (see, e.g., Refs. [65, 114, 115]). However, if a GW signal is present in the data but overlaps with a noise artifact, it is deemed “contaminated” and will be vetoed. In other words, we do not yet have a sufficiently reliable method to separate signals from overlapping, unidentified noise artifacts. Consequently, the false dismissal probability cannot be easily quantified for the vetoes discussed here with the presence of noise artifacts.

1. HMM searches: Merger remnants

a. Known-line veto

The first veto we use in the HMM search involves comparing the frequency path of each signal candidate against all known instrumental lines present in either the Hanford or Livingston detector to see if there is any overlap [149]. We increase the width of the Viterbi path $\delta f \approx 10^{-6} f_i + 8 \delta f_{\text{SFT}}$, where f_i is a given frequency anywhere along the path and δf_{SFT} is the SFT frequency bin width. The first part accounts for the Doppler modulation due to Earth’s sidereal motion, and the second part accounts for the additional data used in the \mathcal{F} -statistic calculation. We consider candidates with a wide range

¹⁰ Note that a factor 2 is missing compared to Ref. [134]. In the present work, these resolutions are used to compute the distance from a template placed on the parameter space rather than the total extent of this template.

of sky positions, start times, and total durations. Because most search timescales are much shorter than a year, we conservatively choose not to factor in Earth’s orbital Doppler modulation in this initial veto. The few longer-duration candidates (with $T_{\text{obs}} > \text{a few months}$) that could have been vetoed if Earth’s orbital Doppler modulation had been included are addressed in later steps.

b. Single-interferometer veto

Next we use a technique that vetoes candidates caused by noise artifacts in a single detector that are not yet well understood or identified in the official release of O4a known instrumental lines. The detailed criteria are as follows: We run the search with each interferometer individually. Then, a signal candidate can be vetoed as an unknown instrumental line if the detection statistic in one detector is below threshold while the other is greater than the detection statistic from the combined detector search, and if the Viterbi paths of the latter two searches overlap.

c. Manual inspection

Because the previous two veto procedures are designed for following up long-duration CW search candidates, they are not always able to effectively identify candidates from a short-duration search. Thus, there are many candidates that still remain at this stage. We visually inspect each candidate using a variety of approaches designed to distinguish signals from noise artifacts, typically applying multiple checks to each candidate for confirmation. These approaches are listed here.

1. We relax certain criteria of the single interferometer veto based on close inspection; for example, we veto any candidate whose detection statistic is below threshold in one detector and anomalously high in the other (e.g., $\bar{\mathcal{L}} \gtrsim 100$), regardless of the detection statistic from the original combined search.
2. We compare the frequency drift of the candidate across T_{obs} against the expected frequency drift of a real signal. If the candidate frequency remains within a single bin for at least three quarters of the total duration, *and* if it is not identified by at least one other configuration with a larger value of T_{coh} (which is by design more sensitive to signals with smaller frequency drifts), then it is unlikely to be a real signal. For example, one candidate from GW231123 has an apparent $\dot{f}_0 < 10^{-8} \text{ Hz s}^{-1}$, but the expected frequency drift optimized for $T_{\text{coh}} = 14.2 \text{ m}$ is $\dot{f}_0 \sim 7 \times 10^{-7} \text{ Hz s}^{-1}$. This veto criterion is typically applied in conjunction with at least one other veto as a cross-check.
3. We assess whether the detection statistic of a given candidate is consistent with expectations for a signal

from the remnant BH given its distance estimate. If the statistic is significantly larger than expected (e.g., $\bar{\mathcal{L}} \gtrsim 100$), we can veto the candidate as a loud but unidentified noise artifact. This veto is applied only in conjunction with at least one other veto. For example, if the detection statistic combining two detectors is significantly larger than expected for the target’s distance, it is usually accompanied by the candidate being much louder in one detector than in the other, in which case the candidate is also vetoed according to the first manual inspection criterion described above.

4. We examine the signal candidate in the spectrogram from each interferometer in the relevant frequency band and time segment to determine whether it overlaps with any visible noise artifacts.
5. Finally, if two candidates with the same frequency path are found using two different search configurations—in other words, if the same candidate is identified in two different searches—and if one of them has been vetoed by any of the above criteria, we consider the other likely to have arisen from the same artifacts. This veto criterion is typically applied in conjunction with other manual inspections as a cross-check.

2. BSD search: Cygnus X-1

For the BSD search, the 27 stage-2 candidates have been filtered by removing outliers associated with known instrumental lines [149] and clustering candidates that shared a common outlier. Six candidates remained after this filtering and are listed in Table VIII. The values reported in the table correspond to the averages of the template parameters between the two detectors, and the quoted uncertainties are obtained by propagating the template resolutions given in Eqs. A10 and A11. As discussed in Sec. III B 2, the other parameters, $\{\text{RA}, \text{Dec}, P, e, \omega\}$, are fixed to the central values listed in Tables I and II, and are therefore identical for all the candidates.

Given the small number of surviving candidates, we conducted the following tests and visually inspected the results for each. Although we did not formally estimate the false dismissal probability, the validity and safety of each veto have been confirmed by applying the follow-up procedure to simulated signals with parameters similar to those of the six candidates. The tests, adapted from previous BSD-based searches (e.g., [65, 114]), are grouped into three categories: (i) consistency of the candidate with its astrophysical parameters, (ii) consistency of the detection statistic with different search configurations, and (iii) a manual inspection where we associated some candidates to artifacts caused by non-Gaussianities in the detector. These tests are described in detail below.

Table VIII. Parameters of the candidates from the BSD search, averaged between the two detectors. For each candidate, we indicate the test by which it was vetoed.

| f_0 [Hz] | a_p [s] | t_{asc} [s] | Veto |
|------------|------------------|-----------------------|--------|
| 32.751 | 38.84 ± 0.82 | $5\,600 \pm 1\,700$ | Stat. |
| 37.131 | 36.09 ± 0.73 | $188\,100 \pm 1\,600$ | Manual |
| 41.176 | 36.79 ± 0.66 | $-4\,383 \pm 1\,400$ | Astro. |
| 41.627 | 38.60 ± 0.65 | $-11\,800 \pm 1\,300$ | Stat. |
| 53.890 | 40.62 ± 0.50 | $204\,700 \pm 950$ | Manual |
| 96.114 | 38.10 ± 0.28 | $223\,790 \pm 570$ | Manual |

a. Known-line veto

Similar to the HMM method, we veto any candidate whose frequency evolution crosses a known instrumental line in either the Hanford or Livingston detector [149]. For a candidate observed at a frequency f_0 , we estimate the modulation range as $f_0 \pm \delta f$, where δf is the modulation size computed with the parameters of the candidate, $\delta f = \max_t \left(\frac{d\Delta\phi(t)}{dt} \right)$, where $\Delta\phi(t)$ is the phase modulation defined in Eq. A2.

b. Clustering

During the candidate selection process, we consider all possible pairs between the outliers of the two detectors. Some redundancy is therefore possible in the list of candidates. Indeed, from the eight candidates passing the known-line veto, two outliers are repeated twice, in association with two different (but similar) outliers in the second detectors. We therefore cluster these two sets of candidates and follow them up conjointly. The number of unique candidates is reduced to six.

c. Astrophysical consistency veto

The second step of the follow-up procedure involves assessing the consistency of each candidate with an astrophysical signal through two key checks. Any candidate failing at least one of these tests is discarded from further analysis.

1. **Frequency refinement.** Since the search demodulates the data using the central frequency of each 1 Hz band, a real signal may not be perfectly corrected if its true frequency differs from this reference. To refine the analysis, we re-run a localized search using the candidate recovered frequency as the reference frequency for demodulation. A genuine signal must persist with similar or improved significance under this correction.

2. **Unmodulated test.** As the search is directed toward a signal emitted from Cygnus X-1, the significance of an astrophysical signal should decrease when using parameters incompatible with the target. In particular, we try to recover the candidate without performing the heterodyne correction. A candidate persisting with comparable significance in this configuration cannot be of astrophysical origin and is therefore vetoed.

d. Statistical consistency check

1. **Sensitivity vetoes.** For a true signal, the detection statistic should scale with the sensitivity of each detector. We check this by normalizing the CR with the median ASD of the 1 Hz band, $\sqrt{S_n}$. We then require

$$\frac{\text{CR}_1}{\sqrt{S_{n_1}}} < 3 \frac{\text{CR}_2}{\sqrt{S_{n_2}}}, \quad (\text{B1})$$

where detector 1 is less sensitive and detector 2 is more sensitive. The factor 3 is a conservative choice and indeed none of the candidate was vetoed by this check.

2. **Cumulative and Uniformity veto.** A genuine astrophysical CW signal should persist throughout the observation run, and we expect: (i) steadily increasing significance as more data are included, and (ii) the candidate to be present in any subset of the data.

We first examined cumulative behavior by computing the CR and signal-to-noise ratio (SNR) over data segments whose durations increased in 30-day steps, for both detectors and for data with and without heterodyne correction. We compared the corrected and uncorrected results, looking for behavior inconsistent with the presence of a signal. One candidate was vetoed because the uncorrected data yielded higher significance for most of the run. A second candidate showed a sharp increase during the first month, indicating the presence of non-stationary noise.

To verify the persistence across subsets, we analyzed one-month segments with various start times, again for both detectors and both data types. Some variation is expected due to varying duty cycles or noise levels, so results were interpreted with tolerance. For the second suspect candidate identified in the cumulative test, corrected and uncorrected data gave similar CR values except in the first month, where the corrected data produced a much higher CR. This confirmed the non-persistent nature of the candidate, and it was vetoed.

e. Manual inspection

Finally, we manually inspect the three remaining candidates. For each, we investigate the spectra of the corrected and uncorrected data using two different frequency resolutions: 5.5×10^{-4} and 1.1×10^{-5} Hz, corresponding to coherence times of 30 min and 1 day, respectively. In all three candidates' bands, we observe strong non-Gaussian noise profiles. Such non-Gaussianities are known to produce artifacts in specific frequency bins during peakmap peak selection [133], and, coincidentally, all the remaining candidates are present in one of these frequency bins. By slightly modifying the resolution of the background estimation used to compute the equalized spectra [133], we force the artifact-affected bins to be moved away from the candidate bins. This modified setup ensures that the candidates' frequency bins were not contaminated by the peakmap creation artifacts. We tested several such configurations, and in all cases the CR of each candidate dropped well below the selection threshold. We also assessed the impact of these modified configurations on simulated signal injections, finding that all injections remained detectable in all configurations. On this basis, we vetoed all remaining candidates.

REFERENCES

- [1] G. Bertone and T. M. P. Tait, *Nature* **562**, 51 (2018).
- [2] G. Bertone and D. Hooper, *Rev. Mod. Phys.* **90**, 045002 (2018).
- [3] E. Oks, *New Astronomy Reviews* **93**, 101632 (2021).
- [4] M. Baryakhtar, L. Rosenberg, and G. Rybka, "Searching for the QCD Dark Matter Axion," (2025), arXiv:2504.10607 [hep-ex].
- [5] R. D. Peccei and H. R. Quinn, *Phys. Rev. Lett.* **38**, 1440 (1977).
- [6] R. D. Peccei and H. R. Quinn, *Phys. Rev. D* **16**, 1791 (1977).
- [7] S. Weinberg, *Phys. Rev. Lett.* **40**, 223 (1978).
- [8] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell, *Phys. Rev. D* **81**, 123530 (2010).
- [9] F. F. Freitas, C. A. Herdeiro, A. P. Morais, A. Onofre, R. Pasechnik, E. Radu, N. Sanchis-Gual, and R. Santos, *Journal of Cosmology and Astroparticle Physics* **2021**, 047 (2021).
- [10] B. Holdom, *Physics Letters B* **166**, 196 (1986).
- [11] M. Goodsell, J. Jaeckel, J. Redondo, and A. Ringwald, *Journal of High Energy Physics* **2009**, 027 (2009).
- [12] J. Jaeckel and A. Ringwald, *Annual Review of Nuclear and Particle Science* **60**, 405 (2010).
- [13] R. Essig *et al.*, in *Snowmass 2013: Snowmass on the Mississippi* (2013) arXiv:1311.0029 [hep-ph].
- [14] L. Hui, J. P. Ostriker, S. Tremaine, and E. Witten, *Phys. Rev. D* **95**, 043541 (2017).
- [15] P. Agrawal, N. Kitajima, M. Reece, T. Sekiguchi, and F. Takahashi, *Physics Letters B* **801**, 135136 (2020).
- [16] M. Fabbrichesi, E. Gabrielli, and G. Lanfranchi, *The physics of the dark photon* (Springer International Publishing, 2021).
- [17] T. Clifton, P. G. Ferreira, A. Padilla, and C. Skordis, *Phys. Rept.* **513**, 1 (2012).
- [18] E. Babichev, L. Marzola, M. Raidal, A. Schmidt-May, F. Urban, H. Veermäe, and M. v. Strauss, *Journal of Cosmology and Astroparticle Physics* **2016**, 016 (2016).
- [19] E. Babichev, L. Marzola, M. Raidal, A. Schmidt-May, F. Urban, H. Veermäe, and M. von Strauss, *Phys. Rev. D* **94**, 084055 (2016).
- [20] K. Aoki and S. Mukohyama, *Phys. Rev. D* **94**, 024001 (2016).
- [21] K. Aoki and K.-i. Maeda, *Phys. Rev. D* **97**, 044002 (2018).
- [22] Y. Manita, K. Aoki, T. Fujita, and S. Mukohyama, *Phys. Rev. D* **107**, 104007 (2023).
- [23] A. Arvanitaki and S. Dubovsky, *Phys. Rev. D* **83**, 044026 (2011).
- [24] H. Yoshino and H. Kodama, *Progress of Theoretical and Experimental Physics* **2014** (2014), 10.1093/ptep/ptu029, 043E02.
- [25] H. Yoshino and H. Kodama, *Progress of Theoretical and Experimental Physics* **2015** (2015), 10.1093/ptep/ptv067, 061E01.
- [26] A. Arvanitaki, M. Baryakhtar, and X. Huang, *Phys. Rev. D* **91**, 084011 (2015).
- [27] A. Arvanitaki, M. Baryakhtar, S. Dimopoulos, S. Dubovsky, and R. Lasenby, *Phys. Rev. D* **95**, 043001 (2017).
- [28] R. Brito, S. Ghosh, E. Barausse, E. Berti, V. Cardoso, I. Dvorkin, A. Klein, and P. Pani, *Phys. Rev. Lett.* **119**, 131101 (2017).
- [29] R. Brito, S. Ghosh, E. Barausse, E. Berti, V. Cardoso, I. Dvorkin, A. Klein, and P. Pani, *Phys. Rev. D* **96**, 064050 (2017).
- [30] M. Baryakhtar, R. Lasenby, and M. Teo, *Phys. Rev. D* **96**, 035019 (2017).
- [31] K. H. M. Chan and O. A. Hannuksela, *Phys. Rev. D* **109**, 023009 (2024).
- [32] V. Cardoso, Óscar J.C. Dias, G. S. Hartnett, M. Middleton, P. Pani, and J. E. Santos, *Journal of Cosmology and Astroparticle Physics* **2018**, 043 (2018).
- [33] D. Baumann, H. S. Chia, and R. A. Porto, *Phys. Rev. D* **99**, 044001 (2019).
- [34] O. A. Hannuksela, K. W. K. Wong, R. Brito, E. Berti, and T. G. F. Li, *Nature Astronomy* **3**, 447 (2019).
- [35] J. Zhang and H. Yang, *Phys. Rev. D* **99**, 064018 (2019).
- [36] W. E. East, *Phys. Rev. D* **96**, 024004 (2017).
- [37] W. E. East and F. Pretorius, *Phys. Rev. Lett.* **119**, 041101 (2017).
- [38] W. E. East, *Phys. Rev. Lett.* **121**, 131104 (2018).
- [39] N. Siemonsen and W. E. East, *Phys. Rev. D* **101**, 024019 (2020).
- [40] R. Brito, V. Cardoso, and P. Pani, *Lect. Notes Phys.* **906**, pp.1 (2020), arXiv:1501.06570 [gr-qc].
- [41] R. Penrose, *Nuovo Cimento Rivista Serie* **1**, 252 (1969).
- [42] W. H. Press and S. A. Teukolsky, *Nature (London)* **238**, 211 (1972).
- [43] Y. B. Zel'Dovich, *Soviet Journal of Experimental and Theoretical Physics Letters* **14**, 180 (1971).
- [44] A. A. Starobinskii, *Soviet Phys JETP* **37**, 28 (1973).
- [45] S. Detweiler, *Phys. Rev. D* **22**, 2323 (1980).
- [46] J. D. Bekenstein, *Phys. Rev. D* **7**, 949 (1973).
- [47] S. R. Dolan, *Phys. Rev. D* **76**, 084001 (2007).

- [48] C. A. Herdeiro, E. Radu, and N. M. Santos, *Physics Letters B* **824**, 136835 (2022).
- [49] M. Isi, L. Sun, R. Brito, and A. Melatos, *Phys. Rev. D* **99**, 084042 (2019).
- [50] M. Baryakhtar, M. Galanis, R. Lasenby, and O. Simon, *Phys. Rev. D* **103**, 095019 (2021).
- [51] N. Siemonsen, C. Mondino, D. Egaña Ugrinovic, J. Huang, M. Baryakhtar, and W. E. East, *Phys. Rev. D* **107**, 075025 (2023).
- [52] T. May, W. E. East, and N. Siemonsen, *Phys. Rev. D* **111**, 044062 (2025).
- [53] J. Aasi *et al.*, *Classical and Quantum Gravity* **32**, 074001 (2015).
- [54] F. Acernese *et al.*, *Classical and Quantum Gravity* **32**, 024001 (2014).
- [55] T. Akutsu *et al.*, *Progress of Theoretical and Experimental Physics* **2021**, 05A101 (2020), <https://academic.oup.com/ptep/article-pdf/2021/5/05A101/37974994/ptaa125.pdf>.
- [56] W. Jia *et al.*, *Science* **385**, 1318 (2024), <https://www.science.org/doi/pdf/10.1126/science.ad08069>.
- [57] B. P. Abbott *et al.*, *Living Reviews in Relativity* **23**, 3 (2020).
- [58] M. Punturo *et al.*, *Proceedings, 14th workshop on gravitational wave data analysis (GWDAA-14): Rome, Italy, January 26-29, 2010*, *Classical and Quantum Gravity* **27**, 194002 (2010).
- [59] M. Evans, R. X. Adhikari, C. Afle, S. W. Ballmer, S. Biscoveanu, S. Borhanian, D. A. Brown, Y. Chen, R. Eisenstein, A. Gruson, A. Gupta, E. D. Hall, R. Huxford, B. Kamai, R. Kashyap, J. S. Kissel, K. Kuns, P. Landry, A. Lenon, G. Lovelace, L. McCuller, K. K. Y. Ng, A. H. Nitz, J. Read, B. S. Sathyaprakash, D. H. Shoemaker, B. J. J. Slagmolen, J. R. Smith, V. Srivastava, L. Sun, S. Vitale, and R. Weiss, arXiv e-prints (2021), [arXiv:2109.09882](https://arxiv.org/abs/2109.09882) [astro-ph.IM].
- [60] K. K. Y. Ng, S. Vitale, O. A. Hannuksela, and T. G. F. Li, *Phys. Rev. Lett.* **126**, 151102 (2021).
- [61] P. S. Aswathi, W. E. East, N. Siemonsen, L. Sun, and D. Jones, arXiv e-prints (2025), [arXiv:2507.20979](https://arxiv.org/abs/2507.20979) [gr-qc].
- [62] R. Abbott *et al.* (The LIGO Scientific Collaboration, the Virgo Collaboration, and the KAGRA Collaboration), *Phys. Rev. D* **105**, 102001 (2022).
- [63] C. Palomba *et al.*, *Phys. Rev. Lett.* **123**, 171101 (2019).
- [64] V. Dergachev and M. A. Papa, *Phys. Rev. Lett.* **123**, 101101 (2019).
- [65] R. Abbott *et al.* (KAGRA, LIGO Scientific, VIRGO), *Phys. Rev. D* **106**, 042003 (2022).
- [66] S. J. Zhu, M. Baryakhtar, M. A. Papa, D. Tsuna, N. Kawanaka, and H.-B. Eggenstein, *Phys. Rev. D* **102**, 063020 (2020).
- [67] L. Tsukada, T. Callister, A. Matas, and P. Meyers, *Phys. Rev. D* **99**, 103015 (2019).
- [68] C. Yuan, Y. Jiang, and Q.-G. Huang, *Phys. Rev. D* **106**, 023020 (2022).
- [69] L. Tsukada, R. Brito, W. E. East, and N. Siemonsen, *Phys. Rev. D* **103**, 083005 (2021).
- [70] L. Sun, R. Brito, and M. Isi, *Phys. Rev. D* **101**, 063020 (2020).
- [71] S. Collaviti, L. Sun, M. Galanis, and M. Baryakhtar, *Classical and Quantum Gravity* **42**, 025006 (2024).
- [72] S. Vermeulen, P. Relton, H. Grote, V. Raymond, C. Afeldt, F. Bergamin, A. Bisht, M. Brinkmann, K. Danzmann, S. Doravari, V. Kringel, J. Lough, H. Lück, M. Mehmet, N. Mukund Menon, S. Nadji, E. Schreiber, B. Sorazu, K. Strain, and H. Wittel, *Nature* **600**, 424 (2021).
- [73] A. S. Göttel, A. Ejlli, K. Karan, S. M. Vermeulen, L. Aiello, V. Raymond, and H. Grote, *Phys. Rev. Lett.* **133**, 101001 (2024).
- [74] L. Aiello, J. W. Richardson, S. M. Vermeulen, H. Grote, C. Hogan, O. Kwon, and C. Stoughton, *Phys. Rev. Lett.* **128**, 121101 (2022).
- [75] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), *Phys. Rev. D* **105**, 063030 (2022).
- [76] A. G. Abac *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), *Phys. Rev. D* **110**, 042001 (2024).
- [77] Q. Yang, L.-W. Ji, B. Hu, Z.-J. Cao, and R.-G. Cai, *Research in Astronomy and Astrophysics* **18**, 065 (2018).
- [78] S. Choudhary, N. Sanchis-Gual, A. Gupta, J. C. Degollado, S. Bose, and J. A. Font, *Phys. Rev. D* **103**, 044032 (2021).
- [79] D. Jones, L. Sun, N. Siemonsen, W. E. East, S. M. Scott, and K. Wette, *Phys. Rev. D* **108**, 064001 (2023).
- [80] D. Jones, N. Siemonsen, L. Sun, W. E. East, A. L. Miller, K. Wette, and O. J. Piccinni, *Phys. Rev. D* **111**, 063028 (2025).
- [81] The LIGO Scientific Collaboration and the Virgo Collaboration and the KAGRA Collaboration, “GWTC-4.0: An Introduction to Version 4.0 of the Gravitational-Wave Transient Catalog,” (2025), [arXiv:2508.18080](https://arxiv.org/abs/2508.18080) [gr-qc].
- [82] The LIGO Scientific Collaboration and the Virgo Collaboration and the KAGRA Collaboration, “GWTC-4.0: Methods for Identifying and Characterizing Gravitational-wave Transients,” (2025), [arXiv:2508.18081](https://arxiv.org/abs/2508.18081) [gr-qc].
- [83] The LIGO Scientific Collaboration and The Virgo Collaboration and the KAGRA Collaboration, “GWTC-4.0: Updating the Gravitational-Wave Transient Catalog with Observations from the First Part of the Fourth LIGO-Virgo-KAGRA Observing Run,” (2025), [arXiv:2508.18082](https://arxiv.org/abs/2508.18082) [gr-qc].
- [84] O. J. Piccinni, P. Astone, S. D’Antonio, S. Frasca, G. Intini, P. Leaci, S. Mastrogiovanni, A. Miller, C. Palomba, and A. Singhal, *Classical and Quantum Gravity* **36**, 015008 (2018).
- [85] J. C. A. Miller-Jones, A. Bahramian, J. A. Orosz, I. Mandel, L. Gou, T. J. Maccarone, C. J. Neijssel, X. Zhao, J. Ziolkowski, M. J. Reid, P. Uttley, X. Zheng, D.-Y. Byun, R. Dodson, V. Grinberg, T. Jung, J.-S. Kim, B. Marcote, S. Markoff, M. J. Rioja, A. P. Rushton, D. M. Russell, G. R. Sivakoff, A. J. Tetarenko, V. Tudose, and J. Wilms, *Science* **371**, 1046 (2021), <https://www.science.org/doi/pdf/10.1126/science.abb3363>.
- [86] T.-W. Wong, F. Valsecchi, T. Fragos, and V. Kalogera, *The Astrophysical Journal* **747**, 111 (2012).
- [87] J. G. Rosa and S. R. Dolan, *Phys. Rev. D* **85**, 044043 (2012).
- [88] P. Pani, V. Cardoso, L. Gualtieri, E. Berti, and A. Ishibashi, *Phys. Rev. D* **86**, 104017 (2012).
- [89] V. P. Frolov, P. Krtouš, D. Kubizňák, and J. E. Santos, *Phys. Rev. Lett.* **120**, 231103 (2018).
- [90] D. Baumann, H. S. Chia, J. Stout, and L. ter Haar, *JCAP* **12**, 006 (2019).
- [91] S. R. Dolan, *Phys. Rev. D* **98**, 104006 (2018).

- [92] N. Siemonsen, T. May, and W. E. East, *Phys. Rev. D* **107**, 104003 (2023), [arXiv:2211.03845 \[gr-qc\]](#).
- [93] The LIGO Scientific Collaboration and the Virgo Collaboration and the KAGRA Collaboration, “GW231123: A binary black hole merger with total mass 190-265 M_{\odot} ,” (2025), [arXiv:2507.08219 \[astro-ph.HE\]](#).
- [94] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, D. Gerosa, L. C. Stein, L. E. Kidder, and H. P. Pfeiffer, *Phys. Rev. Res.* **1**, 033015 (2019).
- [95] C. Brocksopp, A. Tarasov, V. Lyuty, and P. Roche, *Astronomy and Astrophysics* **343** (1999).
- [96] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), *Phys. Rev. D* **106**, 102008 (2022).
- [97] P. Astone, A. Colla, S. D’Antonio, S. Frasca, and C. Palomba, *Phys. Rev. D* **90**, 042002 (2014).
- [98] L. Gou, J. E. McClintock, M. J. Reid, J. A. Orosz, J. F. Steiner, R. Narayan, J. Xiang, R. A. Remillard, K. A. Arnaud, and S. W. Davis, *The Astrophysical Journal* **742**, 85 (2011).
- [99] X. Zhao, L. Gou, Y. Dong, X. Zheng, J. F. Steiner, J. C. A. Miller-Jones, A. Bahramian, J. A. Orosz, and Y. Feng, *The Astrophysical Journal* **908**, 117 (2021).
- [100] L. Gou, J. E. McClintock, R. A. Remillard, J. F. Steiner, M. J. Reid, J. A. Orosz, R. Narayan, M. Hanke, and J. García, *The Astrophysical Journal* **790**, 29 (2014).
- [101] M. Axelsson, R. P. Church, M. B. Davies, A. J. Levan, and F. Ryde, *Monthly Notices of the Royal Astronomical Society* **412**, 2260 (2011), <https://academic.oup.com/mnras/article-pdf/412/4/2260/3334647/mnras0412-2260.pdf>.
- [102] D. J. Walton, J. A. Tomsick, K. K. Madsen, V. Grinberg, D. Barret, S. E. Boggs, F. E. Christensen, M. Clavel, W. W. Craig, A. C. Fabian, F. Fuerst, C. J. Hailey, F. A. Harrison, J. M. Miller, M. L. Parker, F. Rahoui, D. Stern, L. Tao, J. Wilms, and W. Zhang, *The Astrophysical Journal* **826**, 87 (2016).
- [103] R. Duro, T. Dauser, V. Grinberg, I. Miškovíčová, J. Rodríguez, J. Tomsick, M. Hanke, K. Pottschmidt, M. A. Nowak, S. Kreykenbohm, M. Cadolle Bel, A. Bodaghee, A. Lohfink, C. S. Reynolds, E. Kendziorra, M. G. F. Kirsch, R. Staubert, and J. Wilms, *Astronomy and Astrophysics* **589**, A14 (2016), [arXiv:1602.08756 \[astro-ph.HE\]](#).
- [104] A. A. Zdziarski, S. Banerjee, S. Chand, G. Dewangan, R. Misra, M. Szanecki, and A. Niedźwiecki, *The Astrophysical Journal* **962**, 101 (2024).
- [105] T. Kawano, C. Done, S. Yamada, H. Takahashi, M. Axelsson, and Y. Fukazawa, *Publications of the Astronomical Society of Japan* **69**, 36 (2017), https://academic.oup.com/pasj/article-pdf/69/2/36/54678093/pasj_69_2_36.pdf.
- [106] J. A. Tomsick, M. A. Nowak, M. Parker, J. M. Miller, A. C. Fabian, F. A. Harrison, M. Bachetti, D. Barret, S. E. Boggs, F. E. Christensen, W. W. Craig, K. Forster, F. Fürst, B. W. Grefenstette, C. J. Hailey, A. L. King, K. K. Madsen, L. Natalucci, K. Pottschmidt, R. R. Ross, D. Stern, D. J. Walton, J. Wilms, and W. W. Zhang, *The Astrophysical Journal* **780**, 78 (2013).
- [107] H. Krawczynski and B. Beheshtipour, *The Astrophysical Journal* **934**, 4 (2022).
- [108] H. Krawczynski, *General Relativity and Gravitation* **50**, 100 (2018), [arXiv:1806.10347 \[astro-ph.HE\]](#).
- [109] J. M. Miller, C. S. Reynolds, A. C. Fabian, G. Miniutti, and L. C. Gallo, *The Astrophysical Journal* **697**, 900 (2009).
- [110] A. A. Zdziarski, S. Chand, S. Banerjee, M. Szanecki, A. Janiuk, P. Lubiński, A. Niedźwiecki, G. Dewangan, and R. Misra, *The Astrophysical Journal Letters* **967**, L9 (2024).
- [111] L. Sun, A. Melatos, S. Suvorova, W. Moran, and R. J. Evans, *Phys. Rev. D* **97**, 043013 (2018).
- [112] S. Suvorova, L. Sun, A. Melatos, W. Moran, and R. J. Evans, *Phys. Rev. D* **93**, 123009 (2016).
- [113] L. Sun and A. Melatos, *Phys. Rev. D* **99**, 123003 (2019).
- [114] R. Abbott *et al.*, *The Astrophysical Journal* **921**, 80 (2021).
- [115] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **95**, 122003 (2017).
- [116] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **100**, 122002 (2019).
- [117] M. Millhouse, L. Strang, and A. Melatos, *Phys. Rev. D* **102**, 083025 (2020).
- [118] D. Jones and L. Sun, *Physical Review D* **103** (2021).
- [119] D. Beniwal, P. Clearwater, L. Dunn, A. Melatos, and D. Ottaway, *Phys. Rev. D* **103**, 083009 (2021).
- [120] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), *Phys. Rev. D* **105**, 022002 (2022).
- [121] A. M. Knee, H. Du, E. Goetz, J. McIver, J. B. Carlin, L. Sun, L. Dunn, L. Strang, H. Middleton, and A. Melatos, *Phys. Rev. D* **109**, 062008 (2024).
- [122] K. Riles, *Living Reviews in Relativity* **26**, 3 (2023).
- [123] P. Jaranowski, A. Królak, and B. F. Schutz, *Phys. Rev. D* **58**, 063001 (1998).
- [124] C. Cutler and B. F. Schutz, *Physical Review D* **72** (2005), [10.1103/physrevd.72.063006](#).
- [125] LIGO Scientific Collaboration and Virgo Collaboration and KAGRA Collaboration, “LVK Algorithm Library - LALSuite,” Free software (GPL) (2018).
- [126] K. Wette, *SoftwareX* **12**, 100634 (2020).
- [127] A. Viterbi, *IEEE Transactions on Information Theory* **13**, 260 (1967).
- [128] R. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration), *Phys. Rev. D* **106**, 062002 (2022).
- [129] R. Abbott *et al.*, *The Astrophysical Journal Letters* **902**, L21 (2020).
- [130] R. Abbott *et al.*, *The Astrophysical Journal* **935**, 1 (2022).
- [131] O. J. Piccinni, P. Astone, S. D’Antonio, S. Frasca, G. Intini, I. La Rosa, P. Leaci, S. Mastrogiovanni, A. Miller, and C. Palomba, *Phys. Rev. D* **101**, 082004 (2020).
- [132] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **100**, 024004 (2019).
- [133] P. Astone, S. Frasca, and C. Palomba, *Classical and Quantum Gravity* **22**, S1197 (2005).
- [134] P. Leaci and R. Prix, *Phys. Rev. D* **91**, 102003 (2015).
- [135] K. Wette, *Phys. Rev. D* **90**, 122010 (2014).
- [136] E. Capote *et al.*, *Phys. Rev. D* **111**, 062002 (2025).
- [137] S. Soni *et al.*, *Classical and Quantum Gravity* **42**, 085016 (2025).
- [138] D. Ganapathy *et al.* (LIGO O4 Detector Collaboration), *Phys. Rev. X* **13**, 041021 (2023).
- [139] L. Dartez *et al.*, “Characterization of systematic error in Advanced LIGO calibration in the fourth observing run,” (2025).

- [140] M. Wade, J. Betzwieser, D. Bhattacharjee, L. Dartez, E. Goetz, J. Kissel, L. Sun, A. Viets, M. Carney, E. Makelele, and L. Wade, “Toward low-latency, high-fidelity calibration of the LIGO detectors with enhanced monitoring tools,” (2025), [arXiv:2508.08423 \[gr-qc\]](#).
- [141] L. Sun, E. Goetz, J. S. Kissel, J. Betzwieser, S. Karki, A. Viets, M. Wade, D. Bhattacharjee, V. Bossilkov, P. B. Covas, L. E. H. Datrier, R. Gray, S. Kandhasamy, Y. K. Lecoecueche, G. Mendell, T. Mistry, E. Payne, R. L. Savage, A. J. Weinstein, S. Aston, A. Buikema, C. Cahillane, J. C. Driggers, S. E. Dwyer, R. Kumar, and A. Urban, *Classical and Quantum Gravity* **37**, 225008 (2020).
- [142] A. D. Viets, M. Wade, A. L. Urban, S. Kandhasamy, J. Betzwieser, D. A. Brown, J. Burguet-Castell, C. Cahillane, E. Goetz, K. Izumi, S. Karki, J. S. Kissel, G. Mendell, R. L. Savage, X. Siemens, D. Tuyenbayev, and A. J. Weinstein, *Classical and Quantum Gravity* **35**, 095015 (2018).
- [143] A. D. Viets, *Optimizing Advanced LIGO’s scientific output with fast, accurate, clean calibration*, *Phd thesis*, University of Wisconsin-Milwaukee, Milwaukee, WI (2019).
- [144] G. Vajente, Y. Huang, M. Isi, J. C. Driggers, J. S. Kissel, M. J. Szczepańczyk, and S. Vitale, *Phys. Rev. D* **101**, 042003 (2020).
- [145] The LIGO Scientific, Virgo, and KAGRA Collaborations, “Open Data from LIGO, Virgo, and KAGRA through the first part of the fourth observing run,” (2025).
- [146] D. Davis, A. Neunzert, E. Goetz, K. Riles, K. Wette, and M. Lalleman, “Self-gating of O4a h(t) for use in continuous-wave searches,” (2024).
- [147] F. Acernese *et al.*, *Classical and Quantum Gravity* **26**, 204002 (2009).
- [148] P. B. Covas *et al.* (LSC Instrument Authors), *Phys. Rev. D* **97**, 082002 (2018).
- [149] E. Goetz *et al.*, “O4a lines and combs in found in self-gated C00 cleaned data,” (2024).
- [150] N. L. Johnson, A. W. Kemp, and S. Kotz, *Univariate Discrete Distributions* (John Wiley & Sons, Ltd., 2005).
- [151] M. Di Cesare, *All-sky gravitational wave searches for isolated neutron stars: methods and applications to LIGO-Virgo data*, *Master thesis*, Sapienza Università di Roma (2021).
- [152] B. Behnke, M. A. Papa, and R. Prix, *Phys. Rev. D* **91**, 064007 (2015).
- [153] R. Abbott *et al.*, *Physical Review D* **103** (2021), [10.1103/physrevd.103.064017](#).
- [154] R. Abbott *et al.*, *The Astrophysical Journal* **922**, 71 (2021).
- [155] R. Abbott *et al.*, *Physical Review D* **105** (2022), [10.1103/physrevd.105.082005](#).
- [156] C. Palomba, “On the sensitivity of peakmap-based methods for the search of continuous gravitational wave signals,” (2025).