

## Searching for gravitational-wave signals from precessing black hole binaries with the GstLAL pipeline

Stefano Schmidt<sup>1,2,\*</sup> Sarah Caudill,<sup>3,4</sup> Jolien D. E. Creighton,<sup>5</sup> Ryan Magee,<sup>6</sup> Leo Tsukada,<sup>7,8</sup> Shomik Adhichary,<sup>7,8</sup> Pratyusava Baral,<sup>5</sup> Amanda Baylor,<sup>5</sup> Kipp Cannon,<sup>9</sup> Bryce Cousins,<sup>10,7,8</sup> Becca Ewing,<sup>7,8</sup> Heather Fong,<sup>11,9,12</sup> Richard N. George,<sup>13</sup> Patrick Godwin,<sup>14,7,8</sup> Chad Hanna,<sup>7,8,15,16</sup> Reiko Harada,<sup>9,12</sup> Yun-Jing Huang,<sup>7,8</sup> Rachael Huxford,<sup>7,8</sup> Prathamesh Joshi,<sup>7,8</sup> James Kennington,<sup>7,8</sup> Soichiro Kuwahara,<sup>9,12</sup> Alvin K. Y. Li,<sup>6</sup> Duncan Meacher,<sup>5</sup> Cody Messick,<sup>17</sup> Soichiro Morisaki,<sup>18,5</sup> Debnandini Mukherjee,<sup>19,20</sup> Wanting Niu,<sup>7,8</sup> Alex Pace,<sup>7,8</sup> Cort Posnansky,<sup>7,8</sup> Anarya Ray,<sup>5</sup> Surabhi Sachdev,<sup>21,5</sup> Shio Sakon,<sup>7,8</sup> Divya Singh,<sup>7,8</sup> Ron Tapia,<sup>7,16</sup> Takuya Tsutsui,<sup>9</sup> Koh Ueno,<sup>9</sup> Aaron Viets,<sup>22</sup> Leslie Wade,<sup>23</sup> and Madeline Wade<sup>23</sup>

<sup>1</sup>*Nikhef, Science Park 105, 1098 XG Amsterdam, The Netherlands*

<sup>2</sup>*Institute for Gravitational and Subatomic Physics (GRASP), Utrecht University, Princetonplein 1, 3584 CC Utrecht, The Netherlands*

<sup>3</sup>*Department of Physics, University of Massachusetts, Dartmouth, Massachusetts 02747, USA*

<sup>4</sup>*Center for Scientific Computing and Data Science Research, University of Massachusetts, Dartmouth, Massachusetts 02747, USA*

<sup>5</sup>*Leonard E. Parker Center for Gravitation, Cosmology, and Astrophysics, University of Wisconsin-Milwaukee, Milwaukee, Wisconsin 53201, USA*

<sup>6</sup>*LIGO Laboratory, California Institute of Technology, Pasadena, California 91125, USA*

<sup>7</sup>*Department of Physics, The Pennsylvania State University, University Park, Pennsylvania 16802, USA*

<sup>8</sup>*Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, Pennsylvania 16802, USA*

<sup>9</sup>*RESCEU, The University of Tokyo, Tokyo 113-0033, Japan*

<sup>10</sup>*Department of Physics, University of Illinois, Urbana, Illinois 61801, USA*

<sup>11</sup>*Department of Physics and Astronomy, University of British Columbia, Vancouver, British Columbia V6T 1Z4, Canada*

<sup>12</sup>*Graduate School of Science, The University of Tokyo, Tokyo 113-0033, Japan*

<sup>13</sup>*Center for Gravitational Physics, University of Texas at Austin, Austin, Texas 78712, USA*

<sup>14</sup>*LIGO Laboratory, California Institute of Technology, MS 100-36, Pasadena, California 91125, USA*

<sup>15</sup>*Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, Pennsylvania 16802, USA*

<sup>16</sup>*Institute for Computational and Data Sciences, The Pennsylvania State University, University Park, Pennsylvania 16802, USA*

<sup>17</sup>*MIT Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

<sup>18</sup>*Institute for Cosmic Ray Research, The University of Tokyo, 5-1-5 Kashiwanoha, Kashiwa, Chiba 277-8582, Japan*

<sup>19</sup>*NASA Marshall Space Flight Center, Huntsville, Alabama 35811, USA*

<sup>20</sup>*Center for Space Plasma and Aeronomic Research, University of Alabama in Huntsville, Huntsville, Alabama 35899, USA*

<sup>21</sup>*School of Physics, Georgia Institute of Technology, Atlanta, Georgia 30332, USA*

<sup>22</sup>*Concordia University Wisconsin, Mequon, Wisconsin 53097, USA*

<sup>23</sup>*Department of Physics, Hayes Hall, Kenyon College, Gambier, Ohio 43022, USA*



(Received 1 April 2024; accepted 2 July 2024; published 29 July 2024)

Precession in binary black holes (BBH) is caused by the failure of the black hole spins to be aligned and its study can open up new perspectives in gravitational wave astronomy, providing, among other advancements, a precise measure of distance and an accurate characterization of the BBH spins. However, detecting precessing signals is a highly nontrivial task, as standard matched filtering pipelines for gravitational wave searches are built on many assumptions that do not hold in the precessing case. This work details the upgrades made to the GstLAL pipeline to facilitate the search for precessing BBH signals. The implemented changes in the search statistics and in the signal consistency test are then described in detail. The performance of the upgraded pipeline is evaluated through two extensive searches of precessing

\*Contact author: s.schmidt@uu.nl

signals, targeting two different regions in the mass space, and the consistency of the results is examined. Additionally, the benefits of the upgrades are assessed by comparing the sensitive volume of the precessing searches with two corresponding traditional aligned-spin searches. While no significant sensitivity improvement is observed for precessing binaries with mass ratio  $q \lesssim 6$ , a volume increase of up to 100% is attainable for heavily asymmetric systems with largely misaligned spins. Furthermore, our findings suggest that the primary cause of degraded performance in an aligned-spin search targeting precessing signals is not a poor signal-to-noise-ratio recovery but rather the failure of the  $\xi^2$  signal-consistency test. Our work paves the way for a large-scale search for precessing signals, which could potentially result in exciting future detections.

DOI: [10.1103/PhysRevD.110.023038](https://doi.org/10.1103/PhysRevD.110.023038)

## I. INTRODUCTION

According to the general relativistic description of a binary black hole (BBH), the two BH spins play an important role in determining the dynamics of the binary. Indeed, their magnitude and mutual interaction can speed up or slow down the gravitational wave (GW) emission and hence the rate of orbital shrinking. Moreover, if at least one of the two spins is misaligned with the orbital angular momentum [1], the general relativistic interaction between the spin and the orbital angular momentum  $\mathbf{L}$  causes the orbital plane to rotate around a constant axis. This phenomenon is called “precession” and the GW signal emitted in this case acquires a more complicated structure [2–5].

Despite the intricacies inherent in the modeling [6–14], the underlying physical concept is clear: as GW emission attains its maximum along the direction  $\hat{\mathbf{L}}$  of the orbital angular momentum, and given that  $\hat{\mathbf{L}}$  evolves over time, an inertial observer will measure a time modulation of the wave amplitude and phase, as the orbital plane points towards and away from the observer. This effect can be observed by current GW detectors.

Since the first detection of a GW from a BBH coalescence in 2015 [15–17], the LIGO [18], Virgo [19], and KAGRA [20] collaboration (LVK) have detected more than 90 GW signals from the first three observing runs [21–24]. The combined observations of these many signals established that a fraction of the existing binaries do not have the two spins aligned with the orbital angular momentum, although the degree of misalignment favors small values [25,26]. On the other hand, very few signals show strong imprints of precession, as quantified by the precessing inspiral spin parameter  $\chi_P$  [27]. Notable exceptions are GW190521 [28], GW191109 [29], and GW200129 [30], whose spin measurements may indicate support for large amount of spin-orbit misalignment. It is worth noting however that analyses on GW200129 are made difficult by data quality issues [31,32], which could possibly undermine any statement about the precessing nature of the system.

The detection of precessing signals is of primary scientific interest, as the measurement of a strongly precessing signal provides an accurate measurement of the binary distance, breaking the degeneracy between

distance and inclination, with direct impact on the accuracy of cosmological parameters inferred using GW observations [33,34]. Moreover, precession measurements enable a better understanding of the spin distribution in the population of BBHs, which directly impacts our understanding of the BBH formation mechanism [35–38]. Indeed, BBH population studies indicate that heavily precessing systems are more likely to form through dynamic assembly rather than through common evolution [39]. However, currently the fraction of BBHs formed from each channel and a detailed understanding of these channels are still under debate and could possibly be illuminated by further detections of precessing signals.

The most recent LVK GW transient catalogs [21–24] are built by performing sophisticated searches, which can be broadly categorized into unmodeled and modeled. Unmodeled searches, such as those of the coherent WaveBurst pipeline [40–43], search for transient signals with minimal assumptions by looking for coherent excess power between two detectors. Unmodeled searches tend to be less sensitive to compact binary coalescence (CBC) signals than modeled searches but can potentially identify signals from a broader range of sources. In this sense, unmodeled searches have a non-negligible sensitivity to precessing signals, even the most extreme ones. On the other hand, modeled searches rely on matched filtering [44] to correlate the detector’s output with a large number of CBC waveform templates. While modeled searches are optimal when searching for known signals in purely Gaussian noise, they rapidly lose sensitivity for signals not included among the templates.

Despite having possibly detected a few precessing signals, the modeled searches deployed to build the modern catalogs [21–24,45–48] search only for “aligned-spin” systems and do not include the effects of precession in the models used for the template waveforms. This choice drastically simplifies the searches, reduces the computational cost, and makes a systematic search of systems over a broad mass range feasible. Clearly, aligned-spin searches do have some sensitivity to precessing systems, especially for mild precession, allowing for the statements about precession made for the BBH population [25,26].

Furthermore, given the low expected spin from binary neutron stars, the effects of precession are not expected to hurt GW search sensitivity significantly. Additionally, for the highest-mass systems (total mass  $\gtrsim 80$  solar masses) detectable by the LVK, very few precession cycles of the waveform will be present so again, it is unlikely that a cumulative precession effect could significantly hurt detectability. However, we expect precession effects to significantly impact the detectability of strongly precessing binary black hole mergers with a few 10s to 100s of cycles in the LIGO-Virgo sensitivity band [49,50]. For instance, as pointed out in [51], an aligned-spin search targeting neutron-star black-hole (NSBH) systems can miss up to  $\sim 60\%$  of highly precessing sources with mass ratio  $q \gtrsim 6$  and largely misaligned spins (i.e.,  $\chi_p > 0.5$ ).

Modeled searches for GWs are routinely performed by various matched-filter-based GW detection pipelines such as GstLAL [52–57], PyCBC [58–61], the Multi-Band Template Analysis [62,63], the Summed Parallel Infinite Impulse Response pipeline [64–66] among others [44,67–69]. Although they differ in their implementation details, they all rely on a multistep algorithm that (i) correlates GW detector data with a large number of CBC waveform templates to find potential candidates (each called a “trigger”) and (ii) rank the triggers according to their false alarm rate (FAR), which determines how often we can expect such a trigger to occur by chance in a noise-only model. Waveform templates are algorithmically placed in template banks, whose nontrivial generation has developed an active field of research [70–79].

The majority of these pipelines heavily rely on the simplicity of aligned-spin systems to obtain a convenient expression for the matched filter and to reduce the number of templates required to cover the space of interest [44,72]. Indeed, while an aligned-spin system is characterized by four intrinsic parameters (i.e., two masses and the  $z$  components of the two spins), a precessing system must be described by eight intrinsic parameters (i.e., two masses and two 3D spin). For both systems, seven additional extrinsic parameters are required to describe the position and orientation of the binary with respect to the observer. For an aligned-spin system, the search statistics can be maximized over all seven extrinsic parameters, reducing the template’s dependency to only the four intrinsic parameters. However, for a precessing system (and/or a system where imprints from higher-order modes are considered), it is only possible to maximize the search statistics over a few extrinsic parameters. This leaves the template dependent not only on the eight intrinsic parameters but also on the orbital inclination angle between the total angular momentum and the line of sight, and a reference phase [79–81]. Therefore, as the number of parameters required to describe a template increases from 4 to 10, the template placement [82,83] becomes harder and delivers larger template banks, with an associated larger search computational cost.

This challenge, together with the need to introduce a modified expression for the matched filter, has limited the state-of-the-art pipelines to search only for aligned-spin systems.

Several approaches have been developed to tackle the challenges of precessing searches as listed above. Seminal works [4,84,85] focused on modeling precession with a set of phenomenological parameters, which were subsequently used to perform a matched-filtering search [86]. In this approach, the templates do not correspond to real physical signals but they are designed to match adequately physical precessing signals. More recently and in a similar fashion, [87] leveraged the so-called two-harmonic approximation of a precessing signal [88] to achieve a substantial sensitivity increase at a moderate computational cost when searching for precessing signals in the NSBH parameter space.

A different approach consists in directly using the physical waveform to filter the data [80,81,89]. This corresponds to a straightforward extension of the traditional matched-filter pipeline and has the benefits of using the real physical templates and, as such, it does not rely on any approximation of the waveforms, being applicable to virtually any type of binary system (including waveforms not predicted by general relativity). Employing this method typically requires a large number of templates and, as such, it poses the challenge of effectively generating large, high-dimensional template banks at a feasible computational cost. In [79], we addressed this challenge and developed a novel template placement method, specifically designed to generate large template banks, with minimal assumptions on the nature of the GW signal. Hence, by cutting the bank generation time, the exploration of the heavily precessing parameter space was made feasible.

Similar to the direct filtering method of [81], in this paper we design a search method for precessing signals using the features of the GstLAL search pipeline. We demonstrate that our method results in an improvement in the pipeline’s sensitivity to systems with large spin-orbit misalignment, reaching up to 100% increase in the search sensitive volume, depending on the region of the parameter space. We stress that with the new upgrades, the GstLAL pipeline is ready to handle not only a search for precessing signals but also a search for a generic BBH signals, such as those generated by eccentric binaries [90,91] or those with imprints from higher-order modes (HOMs) [92–94].

After briefly reviewing the physics of precessing binaries, in Sec. II we describe our search method and the modifications to the GstLAL pipeline needed to achieve our goal. In Sec. III, after briefly reviewing the template bank generation method [79], we generate and validate two template banks, targeting heavily precessing BBHs. We use these banks to conduct two searches using real data from the LIGO and Virgo interferometers during the third observing run [95,96]. In Sec. IV, we present the results from these searches, validate their performance, and report

the sensitivity improvement in comparison to the aligned-spin version of the pipeline. In Sec. IV, we describe future work which will further improve the pipeline’s sensitivity to precessing binary systems. Final remarks are gathered in Sec. VI.

## II. SEARCH METHOD

The rotation of the orbital plane induced by precession [1] breaks the symmetry of the system under reflection across the orbital plane, characterized by the direction of the orbital angular momentum. This translates into the breaking of the symmetry between the + and the × frequency-domain polarizations  $\tilde{h}_+$ ,  $\tilde{h}_\times$  of the GW generated by an aligned-spin systems which is typically expressed as

$$\tilde{h}_+ \propto i\tilde{h}_\times. \quad (1)$$

As most of the aligned-spin search pipelines use this symmetry to simplify the search problem [44,52,58,62], development of a precessing search pipeline must relax the assumption in Eq. (1) and tackle the general case. In the following discussion, we describe the effects of precession on the GW signals and the consequences for the search method when we do not make any symmetry assumption about the two GW polarizations.

### A. Effects of precession on GW signals and searches

From a qualitative point of view, the effects of precession on a GW signal are more pronounced for heavily spinning asymmetric systems (large mass ratios) observed with an edge-on inclination (i.e., binaries observed with inclination angle  $\iota \simeq \pi/2$ ) [50]. As a binary system has suppressed GW emission on the  $\iota \simeq \pi/2$  direction, even a small variation in the orbital plane orientation translates into a relatively large variation in GW emission. Furthermore, precession is better observable in longer waveforms, primarily because there are more cycles during which precessing effects can accumulate. In typical data analysis applications, where the starting frequency of the signal is usually fixed, longer waveforms correspond to lower total mass. Consequently, with current instruments, the detection of precession is more likely in low-mass, highly asymmetric compact binary systems with very large spins [97,98].

The features of the precessing waveforms are intimately tied to the number of templates needed to search for precession. Indeed, as precession becomes more prominent, for instance at high mass ratios, the number of templates required to cover the parameter space increases significantly compared to aligned-spin searches. Therefore, there is a trade-off between targeting regions where precession has the strongest impact on the signal and the computational cost associated with searching those regions with an adequate number of templates. Similar considerations are crucial for selecting a region within the parameter

space to target and it is essential to focus on areas where sensitivity improvement is more likely to be achieved while keeping a manageable computational cost.

Besides the two BH masses,  $m_1, m_2$ , a binary system is described by two 3D spin vectors  $\mathbf{S}_1$  and  $\mathbf{S}_2$ . It is custom to introduce the dimensionless spin vectors,  $\mathbf{s}_1$  and  $\mathbf{s}_2$ :

$$\mathbf{s}_i = \frac{\mathbf{S}_i}{m_i^2} \quad (2)$$

where each spin is normalized by the maximum allowed spin  $S = m^2$  according to the Kerr solution of the Einstein equation [99]. Thus a generic precessing system is described by eight intrinsic parameters, as opposed to an aligned-spin system, which is described by only four intrinsic parameters, i.e., the two masses and the two L-aligned components of the spins,  $s_{1z}$  and  $s_{2z}$ .

The effect of precession can be quantified by the effective precession spin parameter  $\chi_P$  [27,100], which can reduce the dimensionality of the parameter space. Given a system with mass ratio  $q = \frac{m_1}{m_2} \geq 1$  and spins  $\mathbf{s}_1$  and  $\mathbf{s}_2$ , we define  $s_{1\perp} = |\mathbf{s}_1 \times \hat{\mathbf{L}}|$  as the magnitude of the in-plane component of  $\mathbf{s}_1$  and similarly for  $s_{2\perp}$ . The effective precession spin parameter is equal to

$$\chi_P = \max \left\{ s_{1\perp}, \frac{1}{q} \frac{4 + 3q}{3 + 4q} s_{2\perp} \right\}. \quad (3)$$

$\chi_P$  defines a mapping between the original binary system (with two three-dimensional spins) and a system of reduced dimensionality, where  $\mathbf{s}_1 = (\chi_P, 0, s_{1z})$  and  $\mathbf{s}_2 = (0, 0, s_{2z})$ .<sup>1</sup>

$\chi_P$  is not always able to accurately capture the degree of precession and may prove inadequate for some heavily precessing systems or whenever HOMs are considered [101]. However, despite the introduction of alternative precession measures in the literature [88,97,101,102],  $\chi_P$  retains its status as the most interpretable and widely utilized. Therefore, we will consistently employ  $\chi_P$  throughout this paper to generate a template bank and to summarize our results.

The impact of spin misalignment can be amplified or suppressed by the inclination of the binary system, measured by the angle  $\iota$  between the inclination of the observer’s line of sight and the orbital angular momentum at a specific reference time. When  $\iota$  is close to 0 or  $\pi$ , there is minimal variation in the flux of emitted GWs, resulting in a limited effect of precession on the waveform. Conversely, for  $\iota \simeq \pi/2$ , significant variations in the GW emission occur over time. Therefore, the inclination angle is a crucial parameter to consider when searching for precessing signals.

<sup>1</sup>Two such systems will show similar precessing effects, since by construction the average value of the precessing frequency is roughly the same.




FIG. 1. Effect of precession on a BBH signal. On the top panel we show the GW signal from an aligned-spin BBH, while in the bottom panel we show the effect of precession on the same system. On the y axis we report the observed strain (with arbitrary units). The signal considered is characterized by  $m_1 = 40M_\odot$  and  $m_2 = 3M_\odot$  and  $s_{1z} = -0.2$  and  $s_{2z} = 0.7$ . The precessing signal is chosen with  $\chi_P = 0.8$ , with precession placed only on the first BH spin. The system is in an “edge on” configuration, with inclination  $\iota = \pi/2$ .

In Fig. 1, we show how precession changes the GW signal generated by a BBH. The upper panel reports an aligned-spin BBH, while the bottom panel shows the same signal with a nonzero  $\chi_P$ . We can see how precession changes the amplitude and phase evolution of the waveform. This work is concerned with signals such as the one depicted in the bottom panel of Fig. 1.

### B. Matched filtering

The core of a matched-filtering pipeline relies on a frequentist test of hypotheses, to consistently accept or reject the hypothesis of a signal being in the data  $d$  of a single interferometer. More formally, we compare the “signal” and the “noise” hypotheses through the log likelihood ratio  $\Lambda$  of the data under the two hypotheses:

$$\Lambda = \log \frac{p(d|\text{signal})}{p(d|\text{noise})}. \quad (4)$$

In the “noise” hypothesis the data consist only of Gaussian noise  $n(t)$ , characterized by a known power spectral density (PSD)  $S_n(f)$ . In this case [103,104],

$$p(d|\text{noise}) \propto e^{\frac{1}{2}\langle d|d \rangle}. \quad (5)$$

In Eq. (5),  $\langle \cdot | \cdot \rangle$  denotes the real part of a complex scalar product  $\langle \cdot | \cdot \rangle$  between two time series  $a(t)$ ,  $b(t)$  defined as

$$\langle a|b \rangle = 4 \int_{f_{\min}}^{f_{\max}} df \frac{\tilde{a}^*(f)\tilde{b}(f)}{S_n(f)}, \quad (6)$$

where the integral extends in a suitable frequency range  $[f_{\min}, f_{\max}]$ ,  $\tilde{\cdot}$  denotes the Fourier transform, and  $*$  denotes complex conjugation.

In the “signal” hypothesis, the data consists of a superposition of Gaussian noise  $n(t)$  and a signal  $h(t)$ :

$d(t) = n(t) + h(t)$ . The probability  $p(s|d)$  that a signal is present in some observed data, amounts to the test that the residual  $d - h$  is consistent with Gaussian noise:

$$p(s|d) \propto e^{\frac{1}{2}\langle d-h|d-h \rangle}. \quad (7)$$

For a ground-based detector, the signal  $h(t)$  depends on multiple physical parameters  $\theta$ , which characterize both the source (intrinsic parameters) and the position of the source with respect to the observer (extrinsic parameters):

$$h(t; \theta) = F_+(\delta, \alpha, \Psi)h_+(t; \theta) + F_\times(\delta, \alpha, \Psi)h_\times(t; \theta). \quad (8)$$

The functions  $F_+$ ,  $F_\times$ , also called antenna pattern functions [105,106], denote the interferometer response to the two polarizations of a GW  $h_+$ ,  $h_\times$ , and they depend on the source’s sky location given by right ascension  $\alpha$  and declination  $\delta$  and on the polarization angle  $\Psi$ . For a compact circular binary system,  $h_+$  and  $h_\times$  depend on the 8 intrinsic parameters (i.e., BH masses and spins), the inclination angle  $\iota$ , the reference phase  $\varphi$ , the luminosity distance of the source  $D_L$  and the time of coalescence  $t_c$  [107].

Combining Eqs. (4), (5), and (7), we find the likelihood ratio  $\Lambda$  is given by

$$\Lambda(\theta) = \langle d|\hat{h} \rangle, \quad (9)$$

where  $\hat{h} = \frac{h}{(h|h)^{1/2}}$ . The goal of a matched-filter search is then to maximize at every given time  $\Lambda(\theta)$ , also known as search statistics, as a function of both the intrinsic and extrinsic parameters  $\theta$  of the model. The maximization problem  $\max_\theta \Lambda(\theta)$  can be analytically solved for certain parameters, while a brute force approach is necessary for others.

### 1. The aligned-spin case

As shown in Eq. (1), in the case of aligned-spin circular orbit systems the two frequency-domain polarizations  $\tilde{h}_+$ ,  $\tilde{h}_\times$  are proportional to each other. This can be used to show that the frequency-domain signal model can be written as [44,80]

$$\tilde{h}(f) = \frac{1}{D_{\text{eff}}} e^{i\phi_0} \tilde{h}_+(f; m_1, m_2, s_{1z}, s_{1z}), \quad (10)$$

where  $D_{\text{eff}}$  and  $\phi_0$  absorb the dependence of  $h$  on the extrinsic parameters, including the inclination angle and reference phase.

As Eq. (9) does not depend on an overall amplitude scaling of the template and we can analytically maximize over a constant phase shift, in the aligned-spin case,  $\Lambda(\theta)$  only depends on four intrinsic parameters  $m_1, m_2, s_{1z}, s_{1z}$ , which we need to maximize over by brute force. Maximizing Eq. (9) over the phase  $\phi_0$  in the template given by Eq. (10) gives

$$\max \Lambda^2 = \|\langle s | \hat{h}_+ \rangle\|^2 = (s | \hat{h}_+)^2 + [s | \hat{h}_+]^2 \quad (11)$$

where  $(\cdot | \cdot)$  and  $[\cdot | \cdot]$  denote the real and imaginary part of the scalar product, respectively. The maximisation over the masses and spins is performed by evaluating Eq. (11) for a discrete set of templates, gathered in a template bank. As discussed more in Sec. III, a template bank covers the parameter space of interest with waveform templates so that a random signal within that space will be recovered by at least one template with tolerable loss in the signal-to-noise ratio. For each template in the bank, an aligned-spin matched-filtering pipeline evaluates Eq. (11) as a function of time, by applying a time shift to the data  $s$ . The result of the operation is the so-called SNR time series  $\text{SNR}(t)$ :

$$\text{SNR}(t) = |z(t)|, \quad (12)$$

where  $z(t)$  is the complex SNR time series

$$z(t) = (s | \hat{h}_+)(t) + i[s | \hat{h}_+](t). \quad (13)$$

If we introduce  $\hat{h}_+^{\pi/2}$  as  $\pi/2$  phase shifted version of  $\hat{h}_+$ , such that in frequency domain  $\tilde{h}_+ = i\tilde{h}_+^{\pi/2}$ , then the complex SNR time series becomes [52]

$$z(t) = (s | \hat{h}_+)(t) + i(s | \hat{h}_+^{\pi/2})(t). \quad (14)$$

This is the expression used by the GstLAL pipeline to filter the data. In practice, this amounts to computing the correlation of the data with two orthogonal filters  $\hat{h}_+$  and  $\hat{h}_+^{\pi/2}$  and calculating the SNR by squaring and summing together the two contributions.

### 2. The precessing case

In the case of a precessing signal, Eq. (10) is no longer valid since the two frequency-domain polarizations are not proportional to each other as they were in Eq. (1) for the aligned-spin case. In this case, we are still able to maximize over some of the extrinsic parameters but we lose the ability to analytically maximize over inclination  $\iota$  and reference phase  $\varphi$ , which need to be included in the template bank. As a consequence, a template bank for precessing signals covers a 10-dimensional space, parametrized by the two BH masses, the two 3D spins as well as the angles  $\iota$  and  $\varphi$ .

In this case, the analytical maximization of Eq. (9) yields [27,81,108]

$$\max \Lambda^2 = \frac{(s | \hat{h}_+)^2 + (s | \hat{h}_\times)^2 - 2\hat{h}_{+\times}(s | \hat{h}_\times)(s | \hat{h}_+)}{1 - \hat{h}_{+\times}^2}, \quad (15)$$

where the parameter

$$\hat{h}_{+\times} = (\hat{h}_+ | \hat{h}_\times) \quad (16)$$

controls the discrepancy between the aligned-spin expression for the search statistic in Eq. (11) and the general expression in Eq. (15). Indeed, for  $\hat{h}_{+\times} \rightarrow 0$ , the latter equation reduces to the first, by using  $\tilde{h}_+ \propto i\tilde{h}_\times$ . Note that Eq. (15) is fully generic and does not make any assumption about the nature of the signal. For this reason, besides precession, Eq. (15) is also applicable to circular aligned-spin systems, where imprints of HOMs are considered.

A pipeline designed for detecting precessing and/or HOM signals must incorporate Eq. (15) in the calculation of the SNR time series Eq. (14). This update, combined with a precessing template bank and a suitable signal-consistency test, forms the key components for conducting a precessing search. Neglecting to do so could lead to a substantial decrease in the recovered SNR, thus downgrading any potential discovery of a precessing signal and harming the pipeline sensitivity to precessing signals.

In this work, we compute the SNR time series for a precessing template as

$$z_{\text{prec/HOM}}(t) = (s | \hat{h}_+)(t) + i(s | \hat{h}_\perp)(t) \quad (17)$$

where we introduced the ‘‘orthogonalized’’ template<sup>2</sup>:

$$\hat{h}_\perp = \frac{1}{\sqrt{1 - \hat{h}_{+\times}^2}} (\hat{h}_\times - \hat{h}_{+\times} \hat{h}_+). \quad (18)$$

In the Appendix we show that  $|z_{\text{prec/HOM}}|^2$  is equivalent to the precessing/HOM search statistics Eq. (15). Equation (17) is particularly convenient as it enables the straightforward

<sup>2</sup>Note that, trivially,  $(\hat{h}_\perp | \hat{h}_\perp) = 1$ .

implementation of the search statistics in the GstLAL pipeline by simply modifying the filter waveforms.

It is worth noting that an alternative search statistics, suitable only for precessing signals, is introduced in [80]. While the alternative statistics has the obvious advantage of also maximizing over the reference phase  $\varphi$ , thus reducing the dimensionality of the template bank, it comes at the price of a more complex functional form that translates into an increased computational cost of the search. Moreover, the statistics [80] can be applied in a less general scenario as it does not apply to signals with HOM content. Indeed by implementing the statistics Eq. (17), our pipeline is ready to search for signals with imprints from HOMs, thus reproducing the successful search presented in [94]. See however [109–111] for a promising alternative method to search for systems with HOMs, which does not rely on Eq. (17).

Whenever the value of search statistics passes a given threshold (usually set to 4), the time, the SNR value and the ringing template are recorded, together with some additional quantities: this is called a trigger and it constitutes the first step towards a detection. Clearly, the search statistics itself cannot be used to claim any detection, as it only considers a single detector and may not be adequate in the presence of loud non-Gaussian noise features. Therefore, any detection claim cannot rely solely on selecting the loudest triggers but must instead rely on a more sophisticated ranking statistics [47,54,60,61,111], which takes into account other factors besides the measured values of the search statistics, possibly maximizing or marginalizing over the geometric parameters  $\alpha$ ,  $\delta$ , and  $\Psi$  in the multiple detector case (and not only in the single detector case). As detailed below, the remaining steps of a search pipeline are devoted to build a suitable ranking statistics so as to assess whether a trigger has an astrophysical origin or if it arised by random noise fluctuations.

### C. $\xi^2$ test

The SNR is a very powerful statistics to rank triggers and discerning between GW signal and noise. Nevertheless, in the presence of loud transient noise (glitches), the pipeline may record elevated values of SNR, thus producing false positive detections. To mitigate this effect, the GstLAL pipeline conducts a signal consistency test [52] for each trigger, denoted  $\xi^2$ . Note that it differs from the traditionally utilized  $\chi^2$  signal consistency test [112]. The value of  $\xi^2$ , together with the SNR, is used to identify candidate, leading to a more robust ranking.

The signal consistency test relies on the comparison between the measured complex SNR time series  $z(t)$  and the expected time series  $z_{\text{pred}} = z(0)R(t)$ :

$$\xi^2 = \frac{\int_{-\delta t}^{\delta t} dt |z(t) - z(0)R(t)|^2}{\int_{-\delta t}^{\delta t} dt (2 - 2|R(t)|^2)}, \quad (19)$$

where the integrals extend on a short time window  $[-\delta t, \delta t]$  around the trigger time and  $R(t)$  is a suitable function

describing the template response to a signal matching the template itself.

In the aligned-spin case,  $R(t)$  is given by

$$R(t) = (\hat{h}_+|\hat{h}_+)(t) + i(\hat{h}_+|\hat{h}_+^{\pi/2})(t), \quad (20)$$

which can be easily obtained by plugging Eq. (10) into the complex time series Eq. (14). In the precessing case,  $R(t)$  has a more complicated expression [113]:

$$R_{\text{prec/HOM}}(t) = \frac{1}{2}[(\hat{h}_+|\hat{h}_+)(t) + (\hat{h}_\times|\hat{h}_\times)(t)] + i\frac{1}{2}[(\hat{h}_+|\hat{h}_\times)(t) - (\hat{h}_\times|\hat{h}_+)(t)]. \quad (21)$$

In Sec. IV B, we discuss the performance of the adapted  $\xi^2$  signal consistency test in a real search.

The expression for the template response  $R_{\text{prec/HOM}}(t)$  is only approximate. Indeed, as shown in [113], in the precessing/HOM case the true template response has a more complicated dependence on the two polarizations. This translates into an increased computational cost, and implementing the exact expression would require heavy changes to the pipeline. Luckily, [113] demonstrates that for  $\text{SNR} \lesssim 100$  the approximate template response Eq. (21) is adequate for the purpose of the  $\xi^2$  test, thus keeping the same simple factorization of Eq. (20). Therefore, since a signal has typically a much lower SNR, the use of the approximation is justified.

### D. Event identification stage

For each trigger, the GstLAL pipeline records the instrument  $\{H1, L1, V1\}$ , the trigger time  $t$ , the signal phase  $\phi$  Eq. (10), SNR, the  $\xi^2$  value and the instantaneous sensitivity of each detector, measured in terms of horizon distance  $D$ . Triggers happening on different instruments within a similar time window and with the same template are grouped together, forming a coincident trigger or coincidence.

As discussed above, to discriminate between coincidences arising from random fluctuations and coincidences of astrophysical origin, it is crucial to develop an accurate ranking statistics. Within the GstLAL framework, this is accomplished by computing the ratio between the probability that a coincidence is generated by a signal and the probability that a coincidence is generated by noise only [54,57,114]:

$$\mathcal{L} = \frac{p(\{\vec{D}, \vec{\rho}, \vec{\xi}^2, \vec{t}, \vec{\phi}\}|\text{signal})}{p(\{\vec{D}, \vec{\rho}, \vec{\xi}^2, \vec{t}, \vec{\phi}\}|\text{noise})}, \quad (22)$$

where  $\rho$  is a shorthand for SNR and  $\{\vec{D}, \vec{\rho}, \vec{\xi}^2, \vec{t}, \vec{\phi}\}$  defines a coincidence of one or more instruments. We call  $p(\cdots|\text{signal})$  the signal model, while  $p(\cdots|\text{noise})$  is called

noise model. The statistic Eq. (22) is also known as the likelihood ratio (LR). Both the signal and noise model are parametric probabilistic models whose parameters are set at the time of the search. They are conveniently factorized to exploit the known correlations between the recorded quantities—see also [[57], Eqs. (2) and (8)].

The ranking statistics Eq. (22) is used to assign a FAR to each coincidence. The FAR of a coincident event with LR  $\bar{\mathcal{L}}$  corresponds to the expected rate of background events obtained with a LR  $\mathcal{L} > \bar{\mathcal{L}}$ . A detection is claimed whenever a candidate event reaches a given FAR threshold, usually set to 1/year.

When it comes to precessing searches, the trigger time  $t$ , SNR,  $\xi^2$  and  $D$  keep the same physical meaning they have in the standard case and, for this reason, the terms in the LR involving only such quantities may not require any modification for the precessing case. On the other hand, since Eq. (10) is not valid in the precessing case, the recorded phase  $\phi$  loses its straightforward physical interpretation. As a consequence, the factors of  $\mathcal{L}$  depending on the phase may need to be adapted for a precessing search. The phase dependent terms measure the coherence of the triggers among multiple detectors and in the GstLAL pipeline they correspond to  $p(\vec{\rho}, \vec{t}, \vec{\phi}|\vec{D}, \text{signal})$  and  $p(\vec{t}, \vec{\phi}|\vec{D}, \text{noise})$ . We may refer to them as “coherence” term.

Since we have no reason to believe that the ranking statistics, aside from the “coherence” term, is unsuitable for the precessing case, in this study we maintain without modification the ranking statistics utilized by the pipeline in its latest version [57]. However, the use of “coherence” terms designed for aligned-spin searches in the precessing case could potentially impact the search sensitivity

negatively. Future work may investigate this impact on sensitivity and could develop an improved ranking statistic by computing and implementing an expression for the “coherence” term that is suitable for the precessing case.

### III. TEMPLATE BANKS

To demonstrate the efficacy of the updated GstLAL pipeline, we conduct two runs employing large template banks constructed for precessing signals. One is tailored for mildly asymmetric systems (referred to as the “low  $q$ ” bank), while the other focuses on more asymmetric systems (the “high  $q$ ” bank). For each precessing search, we run an aligned-spin search targeting the same mass range. This comparative approach enables an assessment of the enhanced sensitivity resulting from the use of precessing templates and the modifications in the filtering scheme described in the previous section. The next section Sec. IV presents the results of said comparison. Here, we delve into the template banks’ generation, providing insight into our choices for parameter space and validating their efficacy in covering the target regions. A summary of the template banks generated is presented in Table I.

For the purpose of template placement it is common to define the match  $|\langle \hat{h}_1 | \hat{h}_2 \rangle|$  between two normalized templates  $\hat{h}_1, \hat{h}_2$ , which serves as a measure of similarity between the two templates. As it is standard in the literature, the size and performance of the bank are controlled by the target minimal match  $MM$  [115]: a good bank should have only a small fraction  $\lesssim 10\%$  of random signals with a match lower than  $MM$  with the nearest template of the bank. The actual fraction of signals with a match below  $MM$  depends on the details of the template

TABLE I. We summarize here the most important features of the four template banks considered in this work. We select two regions of the parameter space, labeled “low  $q$ ” and “high  $q$ ,” respectively. For each region, we generate two template banks, one only gathering aligned-spin signals (AS) and another one including precessing signals (P). The aligned-spin banks covers the variable  $M, q, \chi_{\text{eff}}$ , while the precessing template bank cover the variables  $M, q, s_1, \theta_1, s_{2z}, \iota$ . For each template bank we report the range covered by each variable as well as the number of templates (size). The BBH variables not explicitly mentioned in the table are set to 0.

	Mass range	Parameter space	Size	
			AS	P
Low $q$	$m_1, m_2 \in [8, 70]M_\odot$ $q \in [1, 6]$	$s_{1z} \in [-0.99, 0.99]$ $s_{2z} \in [-0.99, 0.99]$ $s_1 \in [0, 0.9]$ $\theta_1 \in [-\pi, \pi]$ $s_{2z} \in [-0.99, 0.99]$ $\iota \in [0, \pi]$	8425	1605625
High $q$	$m_1 \in [15, 70]M_\odot$ $m_2 \in [3, 10]M_\odot$ $q \in [5, 12]$	$s_{1z} \in [-0.99, 0.99]$ $s_{2z} \in [-0.99, 0.99]$ $s_1 \in [0.5, 0.9]$ $\theta_1 \in [-\pi, \pi]$ $s_{2z} \in [-0.99, 0.99]$ $\iota \in [0, \pi]$	27016	2287083



bank generation as well as on the parameter space of choice.

To generate the aligned-spin template banks, we use the state-of-the-art stochastic method [75,76], which builds a template bank through an iterative process. At every step, a new template proposal is randomly sampled and the best match against the other templates of the bank is computed. The proposal is added to the bank only if its best match is lower than the target minimal match. While the method delivers high quality template banks, it requires many expensive match computations. The number of match computation scales with the template bank size and with the dimensionality  $D$  of the templates (i.e., the number of variables required to fully characterize a BBH template). If in the aligned-spin case the number of limited dimension ( $D = 4$ ) allows us to generate a bank at a feasible cost, then the same cannot be said for a precessing template bank, where the larger number of dimensions make the stochastic method too costly.

To avoid these shortcomings, we generate the precessing template banks with our publicly available code MBANK [79,116], specifically designed to cover high dimensional BBH parameter spaces, such as those associated to precession. To place templates, it employs a metric approximation to the distance between templates, which is used to meaningfully define a volume element on the parameter space: the number of templates placed by the algorithm in a given volume  $\mathcal{T}$  is proportional to the volume of  $\mathcal{T}$ , computed by the metric. MBANK implements a novel metric expression, specifically designed to provide an approximation to Eq. (15), rather than to Eq. (11) as done by other metric codes [78,117–121]. Templates are added to the template banks [78,118] by sampling from a suitably trained normalizing flow model [122–124], without any control on their mutual distance. The bank's size grows until a certain fraction  $\eta < 1$  of the parameter space is covered by templates. Our placing method, known in the literature as random, ensures very fast template placement at the price of up to  $\sim 50\%$  larger banks than those produced by the state-of-the-art stochastic approach. While this may seem suboptimal, several studies [82,83,118] argued that in the limit of a large number of dimension, the random method reaches optimality.

As it is common, the nonprecessing template banks sample the variables

$$M, q, s_{1z}, s_{2z},$$

while both the precessing templates bank sample the variables

$$M, q, s_1, \theta_1, s_{2z}, l$$

where  $M = m_1 + m_2$  denotes the total mass of the binary and the first spin is expressed in spherical coordinates:

$$s_x = s \sin \theta, \quad (23)$$

$$s_y = 0, \quad (24)$$

$$s_z = s \cos \theta. \quad (25)$$

In our template bank, precession is only encoded into  $s_x$ ,<sup>3</sup> while all the other in-plane spin components  $s_{1y}, s_{2x}, s_{2y}$  are set to zero.

Note that we set the reference phase  $\varphi = 0$ , despite the fact that the search statistics Eq. (15) does not explicitly maximises over  $\varphi$ . The motivation for this choice is mainly technical: the metric approximation used by MBANK is nearly degenerate if  $\varphi$  is considered, thereby limiting the algorithm's ability to effectively cover the space and resulting in a bank with more templates than needed. This issue was also discussed in [79]. Moreover, we find that omitting the reference phase  $\varphi$  does not significantly impact the bank's ability to recover a generic signal. This is demonstrated by a performance study presented in Fig. 5, where the fitting factor does not depend on the reference phase  $\varphi$ , and in Fig. 4, where the performance of the precessing banks, assessed with injections with a generic reference phase  $\varphi$ , is consistent with that of the aligned-spin banks. This indicates that our choice to neglect  $\varphi$  has minimal effect on the effectualness of the template bank.

To generate the precessing template banks, we use the frequency domain approximant IMRPhenomXP [9], while for the aligned-spin banks we use IMRPhenomD [125]. In all cases, we consider frequencies between  $f_{\min} = 15$  Hz and  $f_{\max} = 1024$  Hz and we use the advanced LIGO O4 design power spectral density (with 190 Mpc range) [126].

To validate a template bank, we select a number of test signals (also called injections) and for each signal  $s(\theta)$ , characterized by  $\theta$ , we compute the fitting factor  $FF$ :

$$FF(\theta) = \max_{\theta' \in \text{bank}} \mathcal{M}(\theta, \theta'), \quad (26)$$

where the match  $\mathcal{M}$  is computed with Eq. (15) evaluated for a normalized signal  $\hat{s}(\theta)$  and a normalized template  $\hat{h}_+(\theta'), \hat{h}_\times(\theta')$ .

The precessing injections are sampled from the relevant mass space using a probability distribution uniform in  $m_1$  and  $q$ . The spins' magnitudes are sampled so that the fourth power of spin is uniformly distributed in  $[0, 1]$ . While this distribution of spins is clearly unphysical, it was chosen to make sure that a large fraction of the injected signals had a large precessing spin: in this way, approximately 70% of the injected signals have  $\chi_P > 0.5$ . The spin directions, the sky location and the binary orientation are all isotropically distributed over the solid angle. The reference phase is uniformly sampled in  $[0, 2\pi]$ . We generate a set of 21947

<sup>3</sup>Here  $s_x$  plays the role of the precession spin parameter  $\chi_P$ .




FIG. 2. Binned templates of the “low  $q$ ” precessing template bank. For each equal size bin, we color code the logarithmic number of templates so that the color is also a measure of the template density. In the different panels, we consider the variables  $M$ ,  $q$ ,  $s_1$ ,  $\theta_1$ ,  $l$ . The bank was first introduced in [79].

precessing injections, with a 30 seconds spacing. This will be used both for validating the banks and for the runs described in Sec. IV. To validate the aligned-spin banks, we use the same 21947 precessing injections as before, with the in-plane spin components set to zero (i.e.,  $s_{1x} = s_{1y} = s_{2x} = s_{2y} = 0$ ).

### A. A template bank for low mass ratio systems

The precessing “low  $q$ ” template bank was introduced in [[79], Sec. V]. It is a large bank covering systems with component masses  $m_1, m_2 \in [8, 70]M_\odot$  with mass ratio limited into  $q \in [1, 6]$ . With a target minimal match of 0.97, the template bank has 1.6 million templates. We also generated an aligned-spin bank in the same mass region, consisting of  $1.4 \times 10^4$  templates. A summary of the two “low  $q$ ” banks is reported in Table I.

In Fig. 2 we plot the template density of the precessing bank, while in Fig. 4 we present the results of an injection study. We note that both the precessing and the aligned-spin banks have an acceptable performance with only  $\sim 10\%$  of the test signals having a match below the 0.97 threshold. The aligned-spin bank recovers  $\sim 15\%$  of the precessing injections with a fitting factor below 0.9. This suggests that the aligned-spin bank approximates a 0.9 target minimal match bank for precessing signal. This remark will be needed to explain some of the search results.

As discussed in [79], the precessing bank provides a very poor coverage in the “low  $q$ , low  $M$ ,” characterized by  $q \lesssim 1.2$  and  $M \lesssim 20M_\odot$ . For this reason, we exclude such region from our searches (and also from our injection studies)<sup>4</sup>: systems in this region are expected to show little or no effect from precession, hence this region is of scarce interest for our purpose of quantifying the benefit of using precessing templates. Finally, we note in Fig. 2 that two small regions at  $\theta = \pm\pi$  and  $q \simeq 5$ , corresponding to the aligned-spin limit, are unphysically populated by many templates. This is a feature of the metric approximation,

<sup>4</sup>To address the issue of poor coverage, an effective strategy could involve using aligned-spin templates to cover the “low  $q$ , low  $M$ ” region.

probably due to the approximant in use, and was also discussed in [79].

### B. A template bank for high mass ratio systems

To target heavily asymmetric precessing systems, we construct the precessing “high  $q$ ” template bank. The bank targets systems with mass ratio  $q \in [5, 12]$ , with primary mass  $m_1 \in [15, 70]M_\odot$  and secondary mass  $m_2 \in [3, 10]M_\odot$ . For the precessing bank, we limit the magnitude of the primary spin  $s_1$  to the range  $[0.5, 0.9]$ . As shown in the fitting factor study Fig. 4, our choice does not impact the bank injection recovery, even in the low  $s_1$  region excluded by the bank. As part of our development, we also generated a bank covering the entire spin range. However, this additional bank, consisting of  $3 \times 10^5$  more templates, did not demonstrate improved coverage of the low  $s_1$  region.

Given the vast size of the parameter space, employing a template bank with a target minimal match of 0.97 would entail dealing with an unwieldy number of templates, reaching as high as 30 million.<sup>5</sup> As modern pipelines are unable to handle banks of similar size, we have pragmatically set a target minimal match requirement of 0.9 for the precessing bank. With the bank now comprising 2.3 million templates, this adjustment helps manage the computational load while still ensuring a reasonable level of template coverage in the search for GW signals.

We also generate an aligned-spin bank, covering the same mass range: with a target minimal match of 0.97, the aligned-spin bank has  $3.9 \times 10^4$  templates. It is noteworthy that the aligned-spin bank is two orders of magnitude smaller than the precessing bank. Consequently in this case, a precessing search is nearly two orders of magnitude more computationally demanding than its aligned-spin counterpart. Other details of the two “high  $q$ ” are reported in Table I. Figure 3 shows the template density of the precessing bank, where we observe that the bulk of the

<sup>5</sup>This number comes from an attempt we made to generate such bank with MBANK. A similar number appears in [87], although referring to a template bank covering a slightly lower mass range.




FIG. 3. Binned templates of the “high  $q$ ” precessing template bank. For each equal size bin, we color code the logarithmic number of templates so that the color is also a measure of the template density. In the different panels, we consider the variables  $M$ ,  $q$ ,  $s_1$ ,  $\theta_1$ ,  $i$ .




FIG. 4. Fitting factor study for the four template banks considered in this work. For each given value of the match, we report the fraction of injections with fitting factor lower than that value. The two aligned-spin banks (AS) and the two precessing banks (P) are tested against the relevant injection set of fully precessing injections used for searching the data. To evaluate their ability to cover the nonprecessing space, the two AS template banks are also tested against a set of aligned-spin injections. The composition of the injection sets is described in the text. In the left panel, we report the results concerning the “low  $q$ ” parameter space, while the right refers to the “high  $q$ ” parameter space. The vertical dashed lines mark the target minimal match of the template banks: 0.9 for the “high  $q$ ” precessing bank and 0.97 for the others.

templates are placed in the high  $q$ , high  $s_1$  region, particularly for systems with  $\theta_1 \simeq \pm\pi/2$ , as seen with an inclination  $i \simeq \pi/2$ . This distribution is consistent with the fact that signals in such regions exhibit the largest precession content. In Fig. 4, we report the results of an injection study for the two banks, showing that both the precessing and “aligned-spin” bank are able to satisfy the relevant minimal match requirements of 0.9 and 0.97, respectively.

In Fig. 5, we study the dependence of the fitting factor of the precessing template bank on the several parameters characterizing a template. We note that the fitting factor is pretty stable across different regions of the parameter space, confirming the quality of our template bank. As the template bank only covers values of  $s_1$  up to 0.9, we expect and observe a drop in accuracy at large values of the primary spin, which lie outside the template bank. The reduced accuracy at high spins seems to be limited only to edge-on systems (i.e.,  $i \simeq \pi/2$ ) with large values of  $\chi_P$  (i.e.,  $\theta \simeq \pm\pi/2$ ). Finally, we see that the coverage is rather

uniform in  $\varphi$ . This justifies our choice of not including the reference phase as part of the variables characterizing the precessing template banks.

For both the precessing and the aligned-spin bank, only  $\sim 10\%$  of the injections are below the target minimal match threshold of 0.9 and 0.97, respectively. This is consistent with the fitting factor results obtained in the “low  $q$ ” case. With only  $\sim 20\%$  of the precessing injections recovered with a match higher than 0.9, the aligned-spin bank struggles to cover the precessing parameter space satisfactorily. Note that in the “low  $q$ ” case, the aligned-spin template bank was able to recover  $\sim 85\%$  of the precessing injections with a match higher than 0.9.

### C. SVD compression

As opposed to the large majority of other pipelines, the GstLAL pipeline performs the matched filtering in time domain. As computing the cross-correlation between two time series in time domain is notoriously more expensive




FIG. 5. Fitting factor of the “high  $q$ ” precessing template bank, as a function of total mass  $M$ , mass ratio  $q$ , tilt  $\theta_1$ , and magnitude  $s_1$  of the first BH, inclination angle  $\iota$  and reference phase  $\varphi$ . For each bin, we report the fifth percentile of the fitting factor distribution.

than in frequency domain,<sup>6</sup> a singular value decomposition (SVD) based template bank compression [52,127,128] scheme has been implemented to mitigate the cost of filtering.

Templates in the banks are grouped together in bins of  $\mathcal{O}(500)$  templates. Each template is then divided in

<sup>6</sup>Indeed, for a time series of  $D$  points, computing the correlation in time domain has a computational cost of  $\mathcal{O}(D^2)$  while in frequency domain the same operation amounts to computing a fast Fourier transform and has a cost of  $\mathcal{O}(D \log D)$ .

different time slices, each with a different sample rate. For each time slice of each bin an SVD decomposition of the templates is performed. The data are then filtered only using a small subset of the SVD basis and the SNR time series for each template is reconstructed by “inverting” the SVD. Filtering the data using the SVD basis results in a reduction of one or two orders of magnitude of the number of filters. Clearly, the bins for the SVD must be carefully chosen: heuristically, each bin should gather very similar waveforms, so that only a small number of SVD basis are required to faithfully represent the templates.

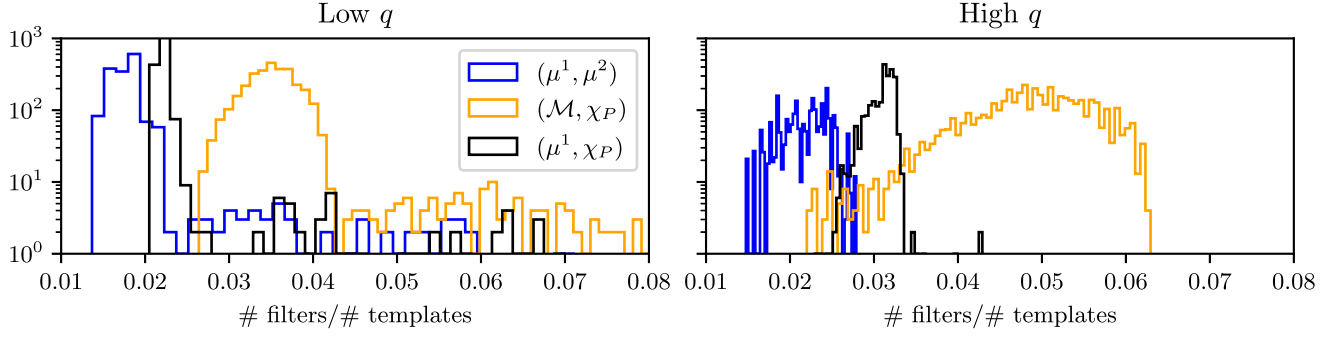


FIG. 6. Compression achieved by the singular value decomposition of the templates of the two precessing template banks considered in this work. For each bin, the compression is measured by the ratio between the number of filters and the number of templates. The number of filters is obtained by averaging, across different time slices, the number of SVD basis. The average is weighted by the sampling rate of each time slice. For different choices of the sorting variables ( $\alpha^1, \alpha^2$ ) used to construct the bank split, we report the histogram of the compression achieved for different bins.

The bins are constructed based on two quantities  $\alpha^1, \alpha^2$ . As a first step, the template bank is split into  $N_{\text{groups}}$  bins based on the  $\alpha^2$  values. Second, groups of  $N_{\text{tmplt}}$  templates are gathered together based on their  $\alpha^1$  value to form an SVD bin. Several choices have been implemented for the two sorting variables ( $\alpha^1, \alpha^2$ ). Common choices for  $\alpha^1$  are the template duration or the chirp mass  $\mathcal{M}$ :

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}. \quad (27)$$

A typical choice for  $\alpha^2$  is the effective spin parameter  $\chi_{\text{eff}} = \frac{m_1 s_{1z} + m_2 s_{2z}}{m_1 + m_2}$  [103]. As in the case of precessing templates,  $\chi_{\text{eff}}$  does not capture the effect of precession, the effective precession spin parameter  $\chi_P$  can be a good option for  $\alpha^2$  instead of  $\chi_{\text{eff}}$ . In [129], the authors suggest to use the PN variables  $\mu^1$  and  $\mu^2$ , first introduced in [130], and they achieve very high compression.

In Fig. 6 we investigate how the compression provided by the SVD varies for the two template banks and for different choices of sorting quantities. The study aims at finding the best sorting quantities for maximal speed up. In our study, we set  $N_{\text{tmplt}} = 500$  and  $N_{\text{groups}} = 20$ .

Our results show that the sorting with variables  $\mu^1$  and  $\mu^2$  introduced in [129] provides the best speed up even in the precessing case. Regardless of the sorting variables considered, the “low  $q$ ” precessing bank shows better compression than the “high  $q$ ” precessing bank. This is expected on the basis that the “high  $q$ ” bank gathers signal with higher precession content and more complex morphology. Therefore, a higher number of SVD basis is needed to faithfully reproduce the waveform.

A comparison with the results in [129] shows that, in the precessing case, the SVD compression is a factor of 2/3 worse than the align spin case. As noted before, this is expected due to the increased complexity of precessing waveforms. The higher number of SVD basis translates

directly into a larger computational cost and it needs to be taken into account when allocating the computational resources for a precessing search.

#### IV. PRECESSING SEARCHES RESULTS

Using the four template banks described in the previous Sec. III, we run four searches on the publicly available LIGO and Virgo data [96] obtained through GWOSC [131], taken during the third observing run between GPS times 1259423400 and 1260081799. As described above, we generate two sets of precessing injections, one for the “low  $q$ ” banks and another for the “high  $q$ ” banks. This choice allows us to directly compare the results of a precessing search with those of an aligned-spin search.

When assessing the sensitivity of a precessing search compared to an aligned-spin search, two competing effects come into play. On one hand, the precessing search enhances the recovered SNR and lowers the  $\xi^2$ , leading to improved sensitivity. On the other hand, the larger bank’s size results in an increased number of background triggers (i.e., false alarms), which can potentially downgrade the significance of any candidate event. Determining whether the improvement in recovered SNR outweighs the increased background triggers, and thus whether a precessing search indeed provides enhanced sensitivity will be the focal point of this section.

In what follows, we evaluate the improvement brought by a precessing search using three different figure of merits: recovered SNR, results from the  $\xi^2$  test and, finally, increase in search sensitivity. To study the sensitivity we used the same injection set introduced for bank validation in Sec. III, with injections sampled uniformly in  $m_1$  and mass ratio inside the mass ranges covered by each template bank.

##### A. Recovered SNR

In Fig. 7, we present a histogram for each injection, illustrating the discrepancy  $\Delta\text{SNR}$  between the SNR

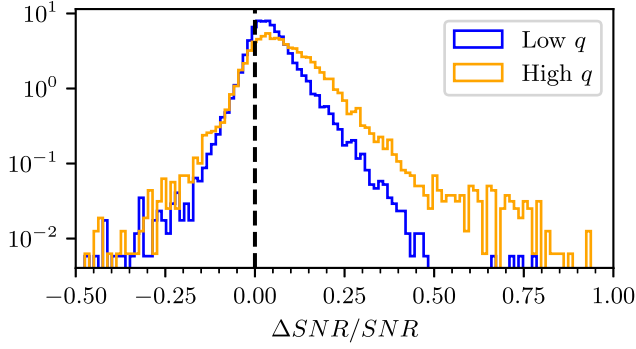


FIG. 7. Histogram with the fractional SNR gain of a precessing search over its aligned-spin counterpart. For the two cases “low  $q$ ” and “high  $q$ ,” we report the discrepancy  $\Delta\text{SNR} = \text{SNR}_{\text{precessing}} - \text{SNR}_{\text{aligned-spin}}$  of the SNR recovered by the precessing and the aligned-spin, normalized by the SNR of each injection. The injections are described in Sec. III and the data refer to the LIGO Livingston detector.




FIG. 8. Histogram with the difference between the  $\xi^2$  measured by a precessing search and its aligned-spin counterpart. For the two cases “low  $q$ ” and “high  $q$ ,” we report the discrepancy  $\Delta\xi^2 = \xi^2_{\text{precessing}} - \xi^2_{\text{aligned-spin}}$ , evaluated on the set of precessing injections described in Sec. III. Data refer to the LIGO Livingston detector.

recovered by the precessing search and the SNR recovered by the aligned-spin search. The SNR difference is divided by the nominal SNR of the injection and we the data refers both to the “low  $q$ ” and “high  $q$ ” cases. This analysis offers a valuable metric to evaluate the improvement brought by filtering the data with precessing templates compared to aligned-spin ones.

A certain degree of scattering of the discrepancies in recovered SNR ( $\Delta\text{SNR}$ ) is expected due to random noise fluctuations. However, our analysis reveals a systematic trend: searches utilizing precessing templates consistently recover a larger fraction of the injected SNR compared to those using aligned-spin templates. This improvement can be substantial, with precessing searches recovering up to 75% more SNR than their aligned-spin counterparts. The systematic increase in recovered SNR fraction with

precessing templates aligns with expectations, given the closer match of precessing templates to injected waveforms shown in Fig. 4. Whether the improved SNR recovery translates into an larger search sensitivity will be assessed in the next sections.

Finally, we observe that the improvement brought by the “high  $q$ ” precessing search over its aligned-spin counterpart is larger than in the “low  $q$ ” case. This is consistent with our results in Fig. 4, which show that the aligned-spin “high  $q$ ” template bank exhibits poorer performance than the “low  $q$ ” one in recovering precessing injections. This is a direct consequence of the effects of precession being more prominent for asymmetric systems.

## B. $\xi^2$ test results

We present in this section results from the  $\xi^2$  consistency test. We will focus on the results of the “high  $q$ ” runs, comparing results from aligned-spin and precessing template banks. First of all, in Fig. 8 we report the difference between the  $\xi^2$  values measured on the injected signals by the precessing and by the aligned-spin search. We note that in most cases, the  $\xi^2$  measured by the precessing pipeline is lower than in the aligned-spin case. This is because the precessing templates are more similar to the injected signals, thus providing a more accurate prediction of the SNR time series and consequently yielding lower  $\xi^2$  values. Moreover, this demonstrates the benefits of the updated  $\xi^2$  expression in Eq. (21) over the standard expression Eq. (20).

In Figs. 9 and 10, we report the values of SNR and  $\xi^2$  for the aligned-spin and precessing search, respectively. The values refer to the LIGO-Livingston interferometer (L1). As clear from the left panel of Fig. 10, for the precessing run, the  $\xi^2$  is able to separate real signals from the background for SNR higher than  $\simeq 9$ . The result is consistent with the standard behavior on a set of aligned-spin injections [56].

In the Figs. 9–10 we see that both runs, aligned-spin and precessing, were able to detect the GW event GW191204\_171526 [24]: the associated trigger is well separated from the background and hence was detected in both cases with a FAR  $< \frac{1}{1000} \text{ yr}^{-1}$ . In the figures, one can also see that the two searches record a handful of background triggers with high SNR and low  $\xi^2$ , with the potential of contaminating the background of the searches. A close followup of these points revealed that such triggers were mostly caused by pairs of glitches happening very close in time with each other. This behavior could possibly emulate the amplitude modulation typical of a precessing signal. It is worth noting that all of those potentially spurious candidates had a very high FAR, as a result of other terms of the ranking Eq. (22) downweighting their significance.

By looking at the right panel of Fig. 10, we note that the injection coincidences with SNR  $\gtrsim 8$  are mostly ranked with a low FAR, meaning that they are mostly recovered. In the region characterized by SNR  $\simeq 10$  and  $\xi^2 \simeq 2$ , the FAR




FIG. 9. “SNR- $\xi^2$ ” plot for the “high  $q$ ” aligned-spin search. On the left panel, we report the SNR and  $\xi^2$  of the background triggers, not associated to any injection, and the injection triggers, corresponding to an injection. On the right panel, we only report the injection triggers, colored by the logarithm of the FAR assigned by the pipeline. Data refer to the LIGO Livingston detector. The red star refers to the trigger produced by the GW event GW191204\_171526 [24], detected by our search with  $\text{FAR} < \frac{1}{1000} \text{ yr}^{-1}$ .




FIG. 10. “SNR- $\xi^2$ ” plot the “high  $q$ ” precessing search. On the left panel, we report the SNR and  $\xi^2$  of the background triggers, not associated to any injection, and the injection triggers, corresponding to an injection. On the right panel, we only report the injection triggers, colored by the logarithm of the FAR assigned by the pipeline. Data refer to the LIGO Livingston detector. The red star refers to the trigger produced by the GW event GW191204\_171526 [24], detected by our search with  $\text{FAR} < \frac{1}{1000} \text{ yr}^{-1}$ .

slightly drops, suggesting that the  $\xi^2$  test loses its efficacy. This can be probably improved by using a denser template bank (at a much larger computational cost), by implementing the unapproximated  $\xi^2$  test [113] or by improving the ranking statistics model.

In Fig. 9, it is evident that, in the case of the aligned-spin bank, the  $\xi^2$  values for injection signals are considerably larger than in the precessing case. Clearly, an aligned-spin template is not able to model accurately the response of a heavily precessing signal, though measuring a high  $\xi^2$  value. Consequently, injection and background triggers are not well separated in the SNR- $\xi^2$  plane and the  $\xi^2$  has dramatically downgraded performance.

It is perhaps surprising that the reduced sensitivity of the aligned-spin search is not caused by a poor SNR recovery but rather by the failure of  $\xi^2$  signal consistency test. In other words, the sensitivity loss does not arise from using the “wrong” aligned-spin templates to filter precessing signals but rather from incorrectly labeling astrophysical triggers as non-Gaussian noise artifacts. This is manifest in the right panel of Fig. 9, where many injections recovered with a

considerably large SNR present very large  $\xi^2$  values, and are consequently assigned a high FAR. This observation can have a deep impact on the planning of future large scale searches for precessing signals, as we will discuss in Sec. V.

An analysis of the “low  $q$ ” runs (not reported here) shows similar conclusions, even though the downranking of injection triggers due to large  $\xi^2$  has less impact on the search’s sensitivity, as will be discussed in the next section.

### C. Sensitivity improvement

We finally measure the improvement in the search sensitivity brought by the use of a precessing bank both in the “low  $q$ ” and in the “high  $q$ ” case. As standard [57,132], the sensitivity of a search is measured by space-time volume (VT), corresponding to the volume reached by a search multiplied by the observing time  $T$ :

$$\text{VT} = T \int_0^\infty dz \epsilon(z, \text{FAR}) \frac{dV}{dz} \frac{1}{1+z}, \quad (28)$$

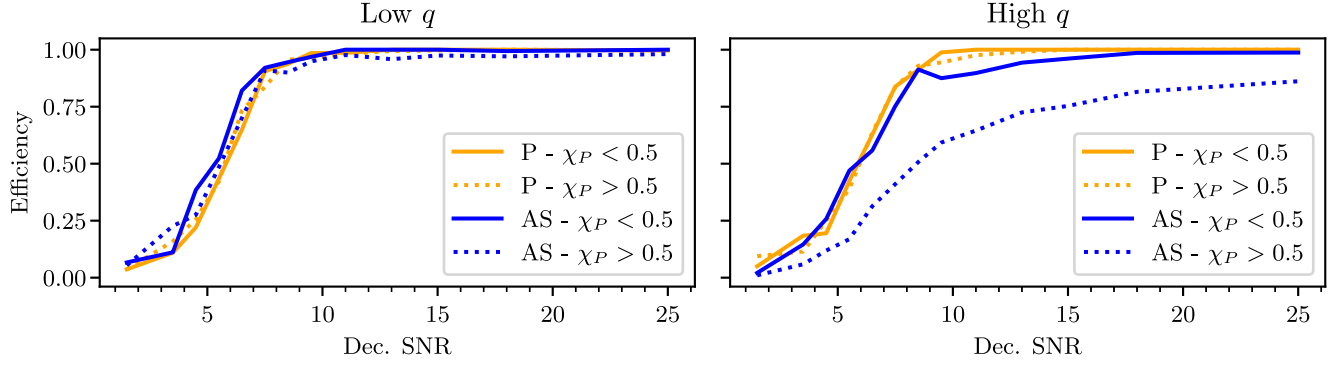


FIG. 11. Efficiency of the BBH searches as a function of the loudest SNR measured by the three detectors, also called the “decisive SNR”. The efficiency is defined as the fraction of signal detected at constant FAR = 1/10 years. In each panel, we show the efficiency pertaining the AS and the P searches, reported separately for  $\chi_P > 0.5$  and  $\chi_P < 0.5$ . The left panel refers to the “low  $q$ ” search, while the right panel to the “high  $q$ ” search.

where  $V(z)$  is the volume of the observable universe up to redshift  $z$  and  $\epsilon(z, \text{FAR})$  is the efficiency of a search, defined as the fraction of injections recovered in an infinitesimal shell between  $z$  and  $z + dz$  with at least a given FAR. Clearly, every estimation for VT heavily depends on the distribution of injection: this motivated the use of the same injection set for the pairs of runs.

In Fig. 11 we report the efficiency (evaluated at a FAR = 1/10 years) of each search as a function “decisive SNR,” defined as the maximum between the three SNRs measured by each detector. In the “low  $q$ ” search, we note that there is not a significant difference between the aligned-spin and the precessing case. Moreover, the efficiency does not depend on the value of  $\chi_P$ . In this case, the improvement in the SNR recovery and in the  $\xi^2$  values brought by a precessing search does not translate into an increased efficiency. In the “high  $q$ ” case, the precessing search has comparable efficacy with the “low  $q$ ” On the

other hand, the aligned-spin search has poorer performance, especially in recovering signals with  $\chi_P > 0.5$ . Thus, in the region covered by the “high  $q$ ” search, the use of precessing templates brought a substantial improvement in the efficiency of the pipeline.

In Fig. 12, we report the ratio  $\frac{VT_P}{VT_{AS}}$  between the VT of by a precessing search and the VT of the corresponding aligned-spin search. The ratio is plotted as a function of FAR and we report different measurement for the low and high  $\chi_P$  regions. We note that for very high values of FAR, the precessing searches always perform better than the aligned-spin ones. This is a direct consequence of the improved SNR recovery offered by the precessing search (see also Sec. IV A and Fig. 7). The picture changes as soon as a lower FAR threshold is considered. Due to the increased background caused by a large number of precessing templates, the “low  $q$ ” precessing search suffers from a sensitivity loss of  $\sim 20\%$  as compared to its aligned-spin

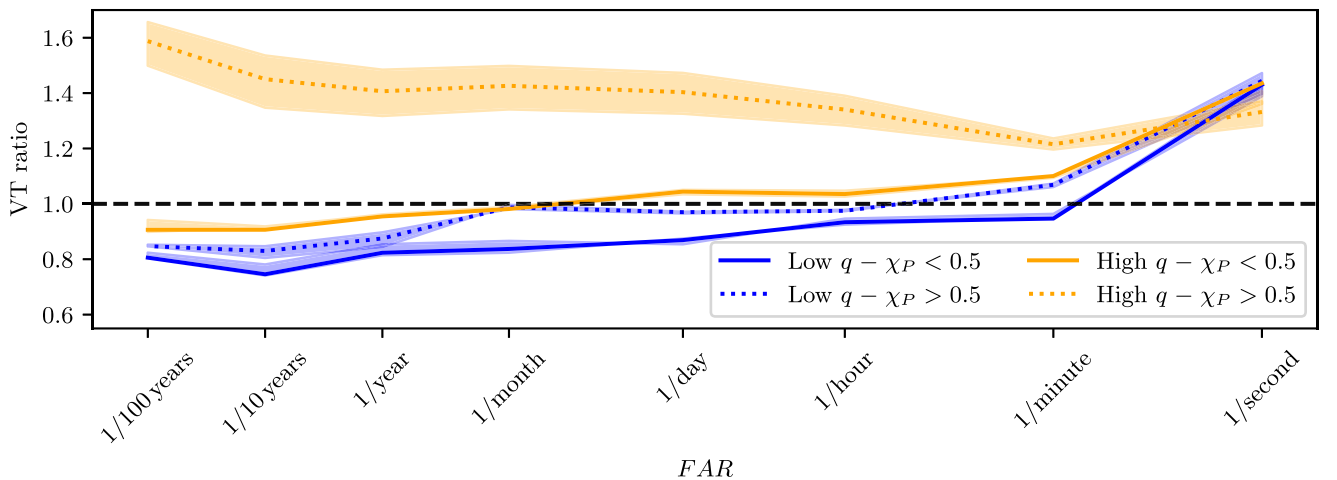


FIG. 12. Sensitive volume increase as a function of the FAR detection threshold. The increase in sensitive volume is computed in terms of the VT ratio, measured between the VT of a precessing search and the VT of an aligned-spin searches. The two blue lines refer to the “low  $q$ ” region, while the orange lines correspond to “high  $q$ .” A dashed line refers to the VT computed only from injections with  $\chi_P > 0.5$ , while a solid line corresponds to  $\chi_P < 0.5$ .






FIG. 13. Sensitive volume increase for different bins in the  $q - \chi_p$  space. The left panel refers to the “low  $q$ ” region while the right panel to the “high  $q$ ” region. For each bin, we compute the VT ratio between a precessing spin search and the corresponding aligned-spin search. The VT ratio is encoded into the colorscale and is computed at a FAR = 1/10 years.

counterpart. Indeed, even for highly precessing signals, with  $\chi_p > 0.5$ , the significance of many triggers is downranked by a large amount of false positives, due to a very large number of templates which can potentially lead to high SNR triggers. A more effective ranking statistics Eq. (22) can mitigate this issue by providing a more robust distinction between signal and noise. This approach was successfully applied in a search targeting precessing NSBH systems [87], where an updated ranking statistic was developed, improving upon its aligned-spin counterpart. Thanks to this enhancement, the authors of [87] achieved up to a  $\sim 50\%$  sensitivity improvement compared to the aligned-spin ranking statistics. Notably, their search did not reveal a decrease in performance for low values of  $\chi_p$ , unlike in our case. Future studies will be needed to understand the limitations of the current ranking statistics and to develop more powerful alternatives that are better suited for detecting precessing signals.

In the “high  $q$ ” scenario, we still observe a decrease in sensitivity for mildly precessing signals  $\chi_p < 0.5$ , even at low FAR. The observation above is still valid: the increased background downranks many potential candidates found by the precessing search. On the other hand, for  $\chi_p > 0.5$  we observe a  $\sim 50\%$  increase of the volume reached by the precessing search. The increase remains stable with the FAR. These considerations are consistent with our observations regarding the measured search efficiency Fig. 11.

It is remarkable, if not surprising, that in the “high  $q$ ” case an improvement in sensitivity was obtained despite the fact that the precessing bank was generated with 0.9 target minimal match. Moreover, the “low  $q$ ” aligned-spin bank, provides an excellent recovery of precessing injections, despite the fact that more than 10% of the injections considered have a match of 0.9 and lower. The two coupled observations seems to indicate that also future searches for precessing signals can employ 0.9 target minimal match

template banks. This has the potential of reducing by orders of magnitude the computational cost as well as the background of the search.

In Fig. 13, the VT ratio is computed for different bins in the  $q - \chi_p$  plane for a fixed FAR of 1/10 years. Results in the left and right panels refer to the “low  $q$ ” and “high  $q$ ,” respectively, and they corroborate what observed above. In the “low  $q$ ” case, we see little variation of the sensitive volume across the space, and for low values of  $q$  and  $\chi_p$ , we observe up to a 40% decrease in sensitivity due to the largely increased background. On the other hand, in the “high  $q$ ” scenario we note that VT ratio can be as high as 100%, depending on the values of  $q$  and  $\chi_p$ . As already noted before, for low  $\chi_p$  the precessing search does not bring any improvement. On the other, for  $\chi_p > 0.5$ , the sensitivity rapidly grows as  $\chi_p$  and  $q$  grow, reaching a 120% gain in the most extreme case.

Our results suggests that targeting more asymmetric systems, such as a NSBH system, could potentially lead to an even larger sensitivity improvement. However as noted in [79], a search in such region struggles with an increased size of the parameter space. Indeed, a template bank in the NSBH region has hundreds of millions of templates and probably a different approach to the search might be needed. In this respect, the work presented in [87] is a promising avenue to target the precessing NSBH parameter space, due its lower computational cost.

## V. OUTLINE OF FUTURE IMPROVEMENTS

Up to this point, we have adapted the GstLAL pipeline to handle precessing templates by modifying the SNR computation routine and making minimal edits to the  $\xi^2$  test, as outlined in Sec. II. However, other components of the pipeline have remained unchanged. Although our

modifications have successfully enhanced the pipeline sensitivity for precessing signals, there might still be room for improvement in the pipeline’s overall performance. Below, we summarize several potential directions for future developments.

First, it is important to acknowledge that the  $\xi^2$  test implemented in this work represents an approximate version of the test proposed in [113]. Implementing the full  $\xi^2$  test would necessitate extensive changes to the pipeline, which are beyond the scope of the present paper. In [113], an injection campaign in both Gaussian and real noise was conducted, suggesting that for signals with  $\text{SNR} \lesssim 20$ , the approximation employed in our implementation of  $\xi^2$  does not introduce significant bias. However, for high SNR signals, the bias introduced by our approximation becomes more significant. We note that the study in [113] did not account for the discreteness of the template bank, which can potentially degrade the performance of the currently implemented  $\xi^2$  test Eq. (21) in a real search scenario. Investigating this aspect is deferred to future work and is essential for determining whether modifying the  $\xi^2$  test can lead to a substantial increase in the performance of a precessing search.

In our searches, we have not employed the “bank  $\xi^2$ ,” introduced in [57]. The “bank  $\xi^2$ ” is an alternative signal consistency test, which checks for a consistent response among several templates within the same SVD bin. It is used as an additional test to the “original”  $\xi^2$  and it has been shown to provide an improvement in the high mass region. Future work should assess the efficacy of the “bank  $\xi^2$ ” applied to precessing signals and possibly deploy it in a real search.

As previously discussed, the ranking statistics given by Eq. (22) may also need adaptation for precessing templates. The key LR term requiring modification is the “coherence” term, as highlighted in Sec. II D. In the precessing case, the measured phase of a signal  $\phi$  loses the straightforward interpretation it has for an aligned-spin signal. Since this information has been crucial in producing the “coherence” term of the ranking statistics, it might not be directly suitable for a precessing template. Deriving an analytical expression that generalizes the “coherence” term to precession is likely to be challenging. In this context, utilizing an unsupervised machine learning model could prove to be an effective solution, allowing for a more flexible and data-driven adaptation of the ranking statistics to precessing templates.

Besides the “coherence” term, other factors of the ranking statistics retain the same physical meaning they have for an aligned-spin search. For this reason, in principle, they do not require modification. However, their predictive power depend on some hyperparameters, which were tuned by only considering aligned-spin searches. Tuning such hyperparameters on a set of precessing injections can offer an increased performance for the search.

We observed in Sec. IV B that the primary loss in sensitivity by the aligned-spin “high  $q$ ” precessing search arises from the degraded performance of the  $\xi^2$  test. This fact can have a profound impact on the search for precessing signals. Indeed, as the SNR recovery is not crucial for performance, it should be possible to deploy a template bank with a very low target minimal match (0.8, or even smaller). The reduced size of the resulting bank will have two favorable implications: (i) it reduces the computational cost and (ii) it reduces the amount of false positives, possibly improving the sensitivity. Of course, a more sparse template bank will also downgrade the performance of the  $\xi^2$  test and future work should explore the trade-off between the benefits of a smaller template bank and the downsides of a less effective signal consistency test. With this strategy, it should become feasible to carry on a precessing search in the NSBH region of the parameter space, where the precession content of a signal is expected to be very strong.

Finally, it is important to stress that the changes we made to the pipeline are fully applicable to systems with HOM content. Similar to the approach taken for precession in this work, future efforts will evaluate the pipeline’s performance in the presence of HOM signals and assess potential sensitivity improvements. Once results from an HOM search become available, the considerations outlined above regarding possible further pipeline upgrades remain relevant.

## VI. FINAL REMARKS

This paper details the modifications made to the GstLAL pipeline to enable a systematic search for GW signals originating from precessing BBH systems. To filter the data, we incorporated the search statistics outlined in [27,81,108] by employing the “orthogonalized” template Eq. (18). This is also discussed in the Appendix. To account for the effects of this modified template response, we made updates to the  $\xi^2$  consistency test. The test now implemented is an approximation to the one proposed in [113].

To assess the effectiveness of our enhancements, we generated two large precessing template banks to search for precessing signals in two distinct regions of the BBH parameter space. These template banks were employed to search one week of publicly available data from the third observing run, spanning between GPS times 1259423400 and 1260081799. The results obtained from these searches were then compared with two identical searches conducted using an aligned-spin template bank, which covered the same mass range.

As discussed in Sec. IV, searching the data with precessing templates results in a slight increase of a signal recovered SNR and to a substantial improvement in the  $\xi^2$  test performance. However, this is achieved at the price of a large increase in the background triggers, due to a template bank being up to two orders of magnitude larger than in the corresponding aligned-spin case. We obtain a substantial improvement of the sensitivity of the pipeline only when

searching strongly precessing systems, characterized by a high mass ratio  $q \gtrsim 5$  and highly precessing spin  $\chi_P \geq 0.5$ . For less strongly precessing systems, we observe a slight decrease of the pipeline sensitivity, as the increased background removes any positive effect of the improved SNR recovery and of the more accurate signal consistency test.

In Sec. V, we outlined a few possible future upgrades of the pipeline, including an updated ranking statistics Eq. (22) and the implementation of the “exact” signal consistency test [113]. Moreover, the pipeline developed for precessing signals can be straightforwardly applied to the search of signals with HOM content. Future work should aim at assessing the sensitivity brought by a similar search.

Our results lead to two significant conclusions that can inform the planning of future precessing searches. The first observation is technical: we find that employing a template bank with a target minimal match of 0.9 is sufficient to achieve a substantial sensitivity improvement for highly precessing signals. As discussed in Sec. IV B, this is due to the fact that the primary sensitivity loss of an aligned-spin search targeting the precessing region is not due to a lack of recovered SNR but rather by a worsening of the  $\xi^2$  value. The observation is crucial for reducing the size of the template bank, consequently lowering computational costs and the associated background in a search. The second observation pertains to the targeted parameter space. Our results suggest that, to attain any sensitivity improvement, a precessing search should focus solely on asymmetric and heavily precessing systems, characterized by  $q \gtrsim 5$  and  $\chi_P \gtrsim 0.5$ . Current aligned-spin searches have limited efficacy in such an extreme region, and indeed, signals in this range have not been detected by past aligned-spin searches.

Our work provides the community with the tools to search for precessing signals. Moreover, our investigations enable the community to select an appropriate parameter space to target with a precessing search, where a large sensitivity improvement, and possibly new detections, are within reach.

### ACKNOWLEDGMENTS

We thank Melissa Lopez Portilla for sharing her knowledge on glitches and detector characterization. S. S. is supported by the research program of the Netherlands Organisation for Scientific Research (NWO). S. C. is supported by the National Science Foundation under Grant No. PHY-2309332. The authors are grateful for computational resources provided by the LIGO Laboratory and supported by National Science Foundation Grants No. PHY-0757058 and No. PHY-0823459. This material is based upon work supported by NSF’s LIGO Laboratory, which is a major facility fully funded by the National Science Foundation. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation (NSF) and operates under cooperative

Agreement No. PHY-1764464. The authors are grateful for computational resources provided by the Pennsylvania State University’s Institute for Computational and Data Sciences (ICDS) and the University of Wisconsin Milwaukee Nemo and support by NSF Grants No. PHY-2011865, No. NSF OAC-2103662, No. NSF PHY-1626190, No. NSF PHY-1700765, No. NSF PHY-2207728, No. NSF PHY-2207594, and No. PHY-2309332 as well as by the research program of the Netherlands Organization for Scientific Research (NWO). This paper carries LIGO Document No. LIGO-P2400044. This research has made use of data or software obtained from the Gravitational Wave Open Science Center [133], a service of LIGO Laboratory, the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA. LIGO Laboratory and Advanced LIGO are funded by the United States National Science Foundation (NSF) as well as 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 GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain. K. A. G. R. A. is supported by Ministry of Education, Culture, Sports, Science and Technology (MEXT), Japan Society for the Promotion of Science (JSPS) in Japan; National Research Foundation (NRF) and Ministry of Science and ICT (MSIT) in Korea; Academia Sinica (AS) and National Science and Technology Council (NSTC) in Taiwan.

### APPENDIX: TEMPLATE ORTHOGONALIZATION

In this Appendix we show that the search statistics Eq. (15) is equivalent to

$$\max \Lambda = (s|\hat{h}_+)^2 + (s|\hat{h}_\perp)^2, \quad (\text{A1})$$

where the “orthogonalized” template  $\hat{h}_\perp$  was introduced in Eq. (18):

$$\hat{h}_\perp = \frac{\hat{h}_\times - \hat{h}_{+\times}\hat{h}_+}{\sqrt{1 - \hat{h}_{+\times}^2}}. \quad (\text{A2})$$

Equation (A1) is more computationally convenient than Eq. (15). Moreover it is well suited to be incorporated into the GstLAL pipeline, since it only involves a minimal change in the template generation routine and no changes in the filtering routine.

By defining  $s_{+/\times} = (s|\hat{h}_{+/\times})$  and  $\hat{h}_{+\times} = (\hat{h}_+|\hat{h}_\times)$ , we trivially obtain the following:

$$\begin{aligned} \max \Lambda &= (s|\hat{h}_+)^2 + \frac{1}{1 - \hat{h}_{+\times}^2} (s|\hat{h}_\times - \hat{h}_{+\times}\hat{h}_+)^2, \\ &= s_+^2 + \frac{\hat{h}_{+\times}}{1 - \hat{h}_{+\times}^2} s_+^2 + \frac{1}{1 - \hat{h}_{+\times}^2} s_\times^2 - 2 \frac{\hat{h}_{+\times}^2}{1 - \hat{h}_{+\times}^2} s_+ s_\times, \\ &= \frac{1}{1 - \hat{h}_{+\times}^2} \{s_+^2 + s_\times^2 - 2\hat{h}_{+\times} s_+ s_\times\}. \end{aligned}$$

In the last line, we can easily recognize Eq. (15). This proves the validity of Eq. (17) for filtering the data with a preprocessing and/or HOM template.

- 
- [1] T. A. Apostolatos, C. Cutler, G. J. Sussman, and K. S. Thorne, Spin induced orbital precession and its modulation of the gravitational wave forms from merging binaries, *Phys. Rev. D* **49**, 6274 (1994).
- [2] L. E. Kidder, C. M. Will, and A. G. Wiseman, Spin effects in the inspiral of coalescing compact binaries, *Phys. Rev. D* **47**, R4183 (1993).
- [3] L. E. Kidder, Coalescing binary systems of compact objects to postNewtonian 5/2 order. 5. Spin effects, *Phys. Rev. D* **52**, 821 (1995).
- [4] A. Buonanno, Y.-b. Chen, and M. Vallisneri, Detecting gravitational waves from precessing binaries of spinning compact objects: Adiabatic limit, *Phys. Rev. D* **67**, 104025 (2003); **74**, 029904(E) (2006).
- [5] M. Campanelli, C. O. Lousto, Y. Zlochower, B. Krishnan, and D. Merritt, Spin flips and precession in black-hole-binary mergers, *Phys. Rev. D* **75**, 064030 (2007).
- [6] E. Racine, A. Buonanno, and L. E. Kidder, Recoil velocity at 2PN order for spinning black hole binaries, *Phys. Rev. D* **80**, 044010 (2009).
- [7] A. Bohe, S. Marsat, G. Faye, and L. Blanchet, Next-to-next-to-leading order spin-orbit effects in the near-zone metric and precession equations of compact binaries, *Classical Quantum Gravity* **30**, 075017 (2013).
- [8] A. Bohé, G. Faye, S. Marsat, and E. K. Porter, Quadratic-in-spin effects in the orbital dynamics and gravitational-wave energy flux of compact binaries at the 3PN order, *Classical Quantum Gravity* **32**, 195010 (2015).
- [9] G. Pratten *et al.*, Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes, *Phys. Rev. D* **103**, 104056 (2021).
- [10] H. Estellés, A. Ramos-Buades, S. Husa, C. García-Quirós, M. Colleoni, L. Haegel, and R. Jaume, Phenomenological time domain model for dominant quadrupole gravitational wave signal of coalescing binary black holes, *Phys. Rev. D* **103**, 124060 (2021).
- [11] S. Akçay, R. Gamba, and S. Bernuzzi, Hybrid post-Newtonian effective-one-body scheme for spin-precessing compact-binary waveforms up to merger, *Phys. Rev. D* **103**, 024014 (2021).
- [12] E. Hamilton, L. London, J. E. Thompson, E. Fauchon-Jones, M. Hannam, C. Kalaghatgi, S. Khan, F. Pannarale, and A. Vano-Vinuales, Model of gravitational waves from precessing black-hole binaries through merger and ring-down, *Phys. Rev. D* **104**, 124027 (2021).
- [13] R. Gamba, S. Akçay, S. Bernuzzi, and J. Williams, Effective-one-body waveforms for precessing coalescing compact binaries with post-Newtonian twist, *Phys. Rev. D* **106**, 024020 (2022).
- [14] A. Ramos-Buades, A. Buonanno, H. Estellés, M. Khalil, D. P. Mihaylov, S. Ossokine, L. Pompili, and M. Shiferaw, Next generation of accurate and efficient multipolar precessing-spin effective-one-body waveforms for binary black holes, *Phys. Rev. D* **108**, 124037 (2023).
- [15] B. P. Abbott *et al.*, GW150914: First results from the search for binary black hole coalescence with Advanced LIGO, *Phys. Rev. D* **93**, 122003 (2016).
- [16] B. P. Abbott *et al.*, Properties of the binary black hole merger GW150914, *Phys. Rev. Lett.* **116**, 241102 (2016).
- [17] B. P. Abbott *et al.*, Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914, *Classical Quantum Gravity* **33**, 134001 (2016).
- [18] J. Aasi *et al.*, Advanced LIGO, *Classical Quantum Gravity* **32**, 074001 (2015).
- [19] F. Acernese *et al.*, Advanced Virgo: A second-generation interferometric gravitational wave detector, *Classical Quantum Gravity* **32**, 024001 (2015).
- [20] T. Akutsu *et al.*, Overview of KAGRA: Detector design and construction history, *Prog. Theor. Exp. Phys.* **2021**, 05A101 (2021).
- [21] B. P. Abbott *et al.*, GWTC-1: A gravitational-wave transient catalog of compact binary mergers observed by LIGO and Virgo during the first and second observing runs, *Phys. Rev. X* **9**, 031040 (2019).
- [22] R. Abbott *et al.*, GWTC-2: Compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run, *Phys. Rev. X* **11**, 021053 (2021).
- [23] R. Abbott *et al.*, GWTC-2.1: Deep extended catalog of compact binary coalescences observed by LIGO and Virgo

- during the first half of the third observing run, *Phys. Rev. D* **109**, 022001 (2024).
- [24] R. Abbott *et al.*, GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run, *Phys. Rev. X* **13**, 041039 (2023).
- [25] R. Abbott *et al.*, Population properties of compact objects from the second LIGO-Virgo gravitational-wave transient catalog, *Astrophys. J. Lett.* **913**, L7 (2021).
- [26] R. Abbott *et al.*, Population of merging compact binaries inferred using gravitational waves through GWTC-3, *Phys. Rev. X* **13**, 011048 (2023).
- [27] P. Schmidt, F. Ohme, and M. Hannam, Towards models of gravitational waveforms from generic binaries II: Modelling precession effects with a single effective precession parameter, *Phys. Rev. D* **91**, 024043 (2015).
- [28] R. Abbott *et al.*, GW190521: A binary black hole merger with a total mass of  $150M_{\odot}$ , *Phys. Rev. Lett.* **125**, 101102 (2020).
- [29] R. C. Zhang, G. Fragione, C. Kimball, and V. Kalogera, On the likely dynamical origin of GW191109 and binary black hole mergers with negative effective spin, *Astrophys. J.* **954**, 23 (2023).
- [30] M. Hannam *et al.*, General-relativistic precession in a black-hole binary, *Nature (London)* **610**, 652 (2022).
- [31] E. Payne, S. Hourihane, J. Golomb, R. Udall, R. Udall, D. Davis, and K. Chatziioannou, Curious case of GW200129: Interplay between spin-precession inference and data-quality issues, *Phys. Rev. D* **106**, 104017 (2022).
- [32] R. Macas, A. Lundgren, and G. Ashton, Revisiting the evidence for precession in GW200129 with machine learning noise mitigation, *Phys. Rev. D* **109**, 062006 (2024).
- [33] S. Vitale and H.-Y. Chen, Measuring the Hubble constant with neutron star black hole mergers, *Phys. Rev. Lett.* **121**, 021303 (2018).
- [34] Q. Yun, W.-B. Han, Q. Hu, and H. Xu, Precessing binary black holes as better dark sirens, *Mon. Not. R. Astron. Soc.* **527**, L60 (2023).
- [35] B. Farr, D. E. Holz, and W. M. Farr, Using spin to understand the formation of LIGO and Virgo's black holes, *Astrophys. J. Lett.* **854**, L9 (2018).
- [36] W. M. Farr, S. Stevenson, M. Coleman Miller, I. Mandel, B. Farr, and A. Vecchio, Distinguishing spin-aligned and isotropic black hole populations with gravitational waves, *Nature (London)* **548**, 426 (2017).
- [37] N. K. Johnson-McDaniel, S. Kulkarni, and A. Gupta, Inferring spin tilts at formation from gravitational wave observations of binary black holes: Interfacing precession-averaged and orbit-averaged spin evolution, *Phys. Rev. D* **106**, 023001 (2022).
- [38] S. Vitale, S. Biscoveanu, and C. Talbot, Spin it as you like: The (lack of a) measurement of the spin tilt distribution with LIGO-Virgo-KAGRA binary black holes, *Astron. Astrophys.* **668**, L2 (2022).
- [39] M. Mapelli, Formation channels of single and binary stellar-mass black holes, in *Handbook of Gravitational Wave Astronomy* (Springer, Singapore, 2021), 10.1007/978-981-15-4702-7\_16-1.
- [40] S. Klimentko, I. Yakushin, A. Mercer, and G. Mitselmakher, Coherent method for detection of gravitational wave bursts, *Classical Quantum Gravity* **25**, 114029 (2008).
- [41] S. Klimentko, G. Vedovato, M. Drago, G. Mazzolo, G. Mitselmakher, C. Pankow, G. Prodi, V. Re, F. Salemi, and I. Yakushin, Localization of gravitational wave sources with networks of advanced detectors, *Phys. Rev. D* **83**, 102001 (2011).
- [42] S. Klimentko *et al.*, Method for detection and reconstruction of gravitational wave transients with networks of advanced detectors, *Phys. Rev. D* **93**, 042004 (2016).
- [43] M. Drago *et al.*, Coherent WaveBurst, a pipeline for unmodeled gravitational-wave data analysis, arXiv:2006.12604.
- [44] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, FINDCHIRP: An algorithm for detection of gravitational waves from inspiraling compact binaries, *Phys. Rev. D* **85**, 122006 (2012).
- [45] A. H. Nitz, S. Kumar, Y.-F. Wang, S. Kastha, S. Wu, M. Schäfer, R. Dhurkunde, and C. D. Capano, 4-OGC: Catalog of gravitational waves from compact binary mergers, *Astrophys. J.* **946**, 59 (2023).
- [46] A. H. Nitz, C. D. Capano, S. Kumar, Y.-F. Wang, S. Kastha, M. Schäfer, R. Dhurkunde, and M. Cabero, 3-OGC: Catalog of gravitational waves from compact-binary mergers, *Astrophys. J.* **922**, 76 (2021).
- [47] S. Olsen, T. Venumadhav, J. Mushkin, J. Roulet, B. Zackay, and M. Zaldarriaga, New binary black hole mergers in the LIGO-Virgo O3a data, *Phys. Rev. D* **106**, 043009 (2022).
- [48] A. K. Mehta, S. Olsen, D. Wadekar, J. Roulet, T. Venumadhav, J. Mushkin, B. Zackay, and M. Zaldarriaga, New binary black hole mergers in the LIGO-Virgo O3b data, arXiv:2311.06061.
- [49] I. W. Harry, A. H. Nitz, D. A. Brown, A. P. Lundgren, E. Ochsner, and D. Keppel, Investigating the effect of precession on searches for neutron-star-black-hole binaries with Advanced LIGO, *Phys. Rev. D* **89**, 024010 (2014).
- [50] J. Calderón Bustillo, P. Laguna, and D. Shoemaker, Detectability of gravitational waves from binary black holes: Impact of precession and higher modes, *Phys. Rev. D* **95**, 104038 (2017).
- [51] R. Dhurkunde and A. H. Nitz, Sensitivity of spin-aligned searches for neutron star-black hole systems using future detectors, *Phys. Rev. D* **106**, 103035 (2022).
- [52] C. Messick *et al.*, Analysis framework for the prompt discovery of compact binary mergers in gravitational-wave data, *Phys. Rev. D* **95**, 042001 (2017).
- [53] S. Sachdev *et al.*, The GstLAL search analysis methods for compact binary mergers in Advanced LIGO's second and advanced Virgo's first observing runs, arXiv:1901.08580.
- [54] C. Hanna *et al.*, Fast evaluation of multidetector consistency for real-time gravitational wave searches, *Phys. Rev. D* **101**, 022003 (2020).
- [55] K. Cannon *et al.*, GstLAL: A software framework for gravitational wave discovery, *SoftwareX* **14**, 100680 (2021).
- [56] B. Ewing *et al.*, Performance of the low-latency GstLAL inspiral search towards LIGO, Virgo, and KAGRA's fourth observing run, *Phys. Rev. D* **109**, 042008 (2024).
- [57] L. Tsukada *et al.*, Improved ranking statistics of the GstLAL inspiral search for compact binary coalescences, *Phys. Rev. D* **108**, 043004 (2023).
- [58] T. Dal Canton *et al.*, Implementing a search for aligned-spin neutron star-black hole systems with advanced ground

- based gravitational wave detectors, *Phys. Rev. D* **90**, 082004 (2014).
- [59] S. A. Usman *et al.*, The PyCBC search for gravitational waves from compact binary coalescence, *Classical Quantum Gravity* **33**, 215004 (2016).
- [60] A. H. Nitz, T. Dent, T. Dal Canton, S. Fairhurst, and D. A. Brown, Detecting binary compact-object mergers with gravitational waves: Understanding and Improving the sensitivity of the PyCBC search, *Astrophys. J.* **849**, 118 (2017).
- [61] G. S. Davies, T. Dent, M. Tápai, I. Harry, C. McIsaac, and A. H. Nitz, Extending the PyCBC search for gravitational waves from compact binary mergers to a global network, *Phys. Rev. D* **102**, 022004 (2020).
- [62] T. Adams, D. Buskulic, V. Germain, G. M. Guidi, F. Marion, M. Montani, B. Mours, F. Piergiovanni, and G. Wang, Low-latency analysis pipeline for compact binary coalescences in the advanced gravitational wave detector era, *Classical Quantum Gravity* **33**, 175012 (2016).
- [63] F. Aubin *et al.*, The MBTA pipeline for detecting compact binary coalescences in the third LIGO–Virgo observing run, *Classical Quantum Gravity* **38**, 095004 (2021).
- [64] J. Luan, S. Hooper, L. Wen, and Y. Chen, Towards low-latency real-time detection of gravitational waves from compact binary coalescences in the era of advanced detectors, *Phys. Rev. D* **85**, 102002 (2012).
- [65] Q. Chu, Low-latency detection and localization of gravitational waves from compact binary coalescences, Ph.D. thesis, Western Australia University, 2017.
- [66] Q. Chu *et al.*, SPIIR online coherent pipeline to search for gravitational waves from compact binary coalescences, *Phys. Rev. D* **105**, 024023 (2022).
- [67] S. Privitera, S. R. P. Mohapatra, P. Ajith, K. Cannon, N. Fotopoulos, M. A. Frei, C. Hanna, A. J. Weinstein, and J. T. Whelan, Improving the sensitivity of a search for coalescing binary black holes with nonprecessing spins in gravitational wave data, *Phys. Rev. D* **89**, 024003 (2014).
- [68] C. Capano, I. Harry, S. Privitera, and A. Buonanno, Implementing a search for gravitational waves from binary black holes with nonprecessing spin, *Phys. Rev. D* **93**, 124007 (2016).
- [69] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, New search pipeline for compact binary mergers: Results for binary black holes in the first observing run of Advanced LIGO, *Phys. Rev. D* **100**, 023011 (2019).
- [70] B. S. Sathyaprakash and S. V. Dhurandhar, Choice of filters for the detection of gravitational waves from coalescing binaries, *Phys. Rev. D* **44**, 3819 (1991).
- [71] S. V. Dhurandhar and B. S. Sathyaprakash, Choice of filters for the detection of gravitational waves from coalescing binaries. 2. Detection in colored noise, *Phys. Rev. D* **49**, 1707 (1994).
- [72] B. J. Owen and B. S. Sathyaprakash, Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement, *Phys. Rev. D* **60**, 022002 (1999).
- [73] S. Babak, R. Balasubramanian, D. Churches, T. Cokelaer, and B. S. Sathyaprakash, A template bank to search for gravitational waves from inspiralling compact binaries. I. Physical models, *Classical Quantum Gravity* **23**, 5477 (2006).
- [74] T. Cokelaer, A template bank to search for gravitational waves from inspiralling compact binaries. II. Phenomenological model, *Classical Quantum Gravity* **24**, 6227 (2007).
- [75] I. W. Harry, B. Allen, and B. S. Sathyaprakash, A stochastic template placement algorithm for gravitational wave data analysis, *Phys. Rev. D* **80**, 104014 (2009).
- [76] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, and A. J. Weinstein, Effectual template bank for the detection of gravitational waves from inspiralling compact binaries with generic spins, *Phys. Rev. D* **89**, 084041 (2014).
- [77] S. Roy, A. S. Sengupta, and P. Ajith, Effectual template banks for upcoming compact binary searches in Advanced-LIGO and Virgo data, *Phys. Rev. D* **99**, 024048 (2019).
- [78] A. Coogan, T. D. P. Edwards, H. S. Chia, R. N. George, K. Freese, C. Messick, C. N. Setzer, C. Weniger, and A. Zimmerman, Efficient gravitational wave template bank generation with differentiable waveforms, *Phys. Rev. D* **106**, 122001 (2022).
- [79] S. Schmidt, B. Gadre, and S. Caudill, Gravitational-wave template banks for novel compact binaries, *Phys. Rev. D* **109**, 042005 (2024).
- [80] I. Harry, S. Privitera, A. Bohé, and A. Buonanno, Searching for gravitational waves from compact binaries with precessing spins, *Phys. Rev. D* **94**, 024012 (2016).
- [81] I. Harry, J. Calderón Bustillo, and A. Nitz, Searching for the full symphony of black hole binary mergers, *Phys. Rev. D* **97**, 023004 (2018).
- [82] B. Allen, Performance of random template banks, *Phys. Rev. D* **105**, 102003 (2022).
- [83] B. Allen, Optimal template banks, *Phys. Rev. D* **104**, 042005 (2021).
- [84] T. A. Apostolatos, Search templates for gravitational waves from precessing, inspiraling binaries, *Phys. Rev. D* **52**, 605 (1995).
- [85] T. A. Apostolatos, Construction of a template family for the detection of gravitational waves from coalescing binaries, *Phys. Rev. D* **54**, 2421 (1996).
- [86] A. Buonanno, Y. Chen, Y. Pan, H. Tagoshi, and M. Vallisneri, Detecting gravitational waves from precessing binaries of spinning compact objects. II. Search implementation for low-mass binaries, *Phys. Rev. D* **72**, 084027 (2005).
- [87] C. McIsaac, C. Hoy, and I. Harry, Search technique to observe precessing compact binary mergers in the advanced detector era, *Phys. Rev. D* **108**, 123016 (2023).
- [88] S. Fairhurst, R. Green, C. Hoy, M. Hannam, and A. Muir, Two-harmonic approximation for gravitational waveforms from precessing binaries, *Phys. Rev. D* **102**, 024055 (2020).
- [89] T. Dal Canton, A. P. Lundgren, and A. B. Nielsen, Impact of precession on aligned-spin searches for neutron-star–black-hole binaries, *Phys. Rev. D* **91**, 062010 (2015).
- [90] A. Ramos-Buades, S. Tiwari, M. Haney, and S. Husa, Impact of eccentricity on the gravitational-wave searches for binary black holes: High mass case, *Phys. Rev. D* **102**, 043005 (2020).

- [91] A. Ramos-Buades, S. Tiwari, M. Haney, and S. Husa, Impact of eccentricity on the gravitational wave searches for binary black holes: High mass case, *Phys. Rev. D* **102**, 043005 (2020).
- [92] L. Blanchet, G. Faye, B. R. Iyer, and S. Sinha, The third post-Newtonian gravitational wave polarisations and associated spherical harmonic modes for inspiralling compact binaries in quasi-circular orbits, *Classical Quantum Gravity* **25**, 165003 (2008); **29**, 239501(E) (2012).
- [93] J. Calderón Bustillo, S. Husa, A. M. Sintes, and M. Pürrer, Impact of gravitational radiation higher order modes on single aligned-spin gravitational wave searches for binary black holes, *Phys. Rev. D* **93**, 084019 (2016).
- [94] K. Chandra, J. Calderón Bustillo, A. Pai, and I. W. Harry, First gravitational-wave search for intermediate-mass black hole mergers with higher-order harmonics, *Phys. Rev. D* **106**, 123003 (2022).
- [95] R. Abbott *et al.*, Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo, *SoftwareX* **13**, 100658 (2021).
- [96] R. Abbott *et al.*, Open data from the third observing run of LIGO, Virgo, KAGRA, and GEO, *Astrophys. J. Suppl. Ser.* **267**, 29 (2023).
- [97] S. Fairhurst, R. Green, M. Hannam, and C. Hoy, When will we observe binary black holes precessing?, *Phys. Rev. D* **102**, 041302 (2020).
- [98] R. Green, C. Hoy, S. Fairhurst, M. Hannam, F. Pannarale, and C. Thomas, Identifying when precession can be measured in gravitational waveforms, *Phys. Rev. D* **103**, 124023 (2021).
- [99] C. Heinicke and F.W. Hehl, Schwarzschild and Kerr solutions of Einstein's field equation—an introduction, *Int. J. Mod. Phys. D* **24**, 1530006 (2014).
- [100] P. Schmidt, M. Hannam, and S. Husa, Towards models of gravitational waveforms from generic binaries: A simple approximate mapping between precessing and non-precessing inspiral signals, *Phys. Rev. D* **86**, 104063 (2012).
- [101] L. M. Thomas, P. Schmidt, and G. Pratten, New effective precession spin for modeling multimodal gravitational waveforms in the strong-field regime, *Phys. Rev. D* **103**, 083022 (2021).
- [102] D. Gerosa, M. Mould, D. Gangardt, P. Schmidt, G. Pratten, and L. M. Thomas, A generalized precession parameter  $\chi_p$  to interpret gravitational-wave data, *Phys. Rev. D* **103**, 064067 (2021).
- [103] M. Maggiore, *Gravitational Waves. Vol. 1: Theory and Experiments*, Oxford Master Series in Physics (Oxford University Press, New York, 2007).
- [104] J. Creighton and W. Anderson, *Gravitational-Wave Physics and Astronomy: An Introduction to Theory, Experiment and Data Analysis* (Wiley, Hoboken, NJ, 2011).
- [105] L. S. Finn and D. F. Chernoff, Observing binary inspiral in gravitational radiation: One interferometer, *Phys. Rev. D* **47**, 2198 (1993).
- [106] P. Jaranowski, A. Krolak, and B. F. Schutz, Data analysis of gravitational—wave signals from spinning neutron stars. 1. The signal and its detection, *Phys. Rev. D* **58**, 063001 (1998).
- [107] B. S. Sathyaprakash and B. F. Schutz, Physics, astrophysics and cosmology with gravitational waves, *Living Rev. Relativity* **12**, 2 (2009).
- [108] C. Capano, Y. Pan, and A. Buonanno, Impact of higher harmonics in searching for gravitational waves from non-spinning binary black holes, *Phys. Rev. D* **89**, 102003 (2014).
- [109] D. Wadekar, J. Roulet, T. Venumadhav, A. K. Mehta, B. Zackay, J. Mushkin, S. Olsen, and M. Zaldarriaga, New black hole mergers in the LIGO-Virgo O3 data from a gravitational wave search including higher-order harmonics, [arXiv:2312.06631](https://arxiv.org/abs/2312.06631).
- [110] D. Wadekar, T. Venumadhav, A. K. Mehta, J. Roulet, S. Olsen, J. Mushkin, B. Zackay, and M. Zaldarriaga, A new approach to template banks of gravitational waves with higher harmonics: Reducing matched-filtering cost by over an order of magnitude, [arXiv:2310.15233](https://arxiv.org/abs/2310.15233).
- [111] D. Wadekar, T. Venumadhav, J. Roulet, A. K. Mehta, B. Zackay, J. Mushkin, and M. Zaldarriaga, A new search pipeline for gravitational waves with higher-order modes using mode-by-mode filtering, [arXiv:2405.17400](https://arxiv.org/abs/2405.17400).
- [112] B. Allen,  $\chi^2$  time-frequency discriminator for gravitational wave detection, *Phys. Rev. D* **71**, 062001 (2005).
- [113] S. Schmidt and S. Caudill, A novel signal-consistency test for gravitational-wave searches of generic black hole binaries, [arXiv:2403.17179](https://arxiv.org/abs/2403.17179) [Phys. Rev. D (to be published)].
- [114] K. Cannon, C. Hanna, and J. Peoples, Likelihood-ratio ranking statistic for compact binary coalescence candidates with rate estimation, [arXiv:1504.04632](https://arxiv.org/abs/1504.04632).
- [115] B. J. Owen, Search templates for gravitational waves from inspiraling binaries: Choice of template spacing, *Phys. Rev. D* **53**, 6749 (1996).
- [116] S. Schmidt, *mbank*—metric bank generation for gravitational waves data analysis, <https://mbank.readthedocs.io/en/latest/>.
- [117] B. J. Owen, Search templates for gravitational waves from inspiraling binaries: Choice of template spacing, *Phys. Rev. D* **53**, 6749 (1996).
- [118] C. Messenger, R. Prix, and M. A. Papa, Random template banks and relaxed lattice coverings, *Phys. Rev. D* **79**, 104017 (2009).
- [119] R. Prix, Template-based searches for gravitational waves: Efficient lattice covering of flat parameter spaces, *Classical Quantum Gravity* **24**, S481 (2007).
- [120] D. A. Brown, I. Harry, A. Lundgren, and A. H. Nitz, Detecting binary neutron star systems with spin in advanced gravitational-wave detectors, *Phys. Rev. D* **86**, 084017 (2012).
- [121] C. Hanna *et al.*, Binary tree approach to template placement for searches for gravitational waves from compact binary mergers, *Phys. Rev. D* **108**, 042003 (2023).
- [122] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, Normalizing flows for probabilistic modeling and inference, *J. Mach. Learn. Res.* **22**, 1 (2019).
- [123] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, Neural spline flows, in *Advances in Neural Information Processing Systems, Vancouver, BC*, edited by H. Wallach,

- H. Laroche, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019), Vol. 32.
- [124] I. Kobyzev, S. J. Prince, and M. A. Brubaker, Normalizing flows: An introduction and review of current methods, *IEEE Trans. Pattern Anal. Mach. Intell.* **43**, 3964 (2021).
- [125] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. Jiménez Forteza, and A. Bohé, Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era, *Phys. Rev. D* **93**, 044007 (2016).
- [126] R. Abbott *et al.*, Noise curves used for simulations in the update of the observing scenarios paper, <https://dcc.ligo.org/LIGO-T2000012/public>.
- [127] K. Cannon, A. Chapman, C. Hanna, D. Keppel, A. C. Searle, and A. J. Weinstein, Singular value decomposition applied to compact binary coalescence gravitational-wave signals, *Phys. Rev. D* **82**, 044025 (2010).
- [128] K. Cannon *et al.*, Toward early-warning detection of gravitational waves from compact binary coalescence, *Astrophys. J.* **748**, 136 (2012).
- [129] S. Sakon *et al.*, Template bank for compact binary mergers in the fourth observing run of Advanced LIGO, Advanced Virgo, and KAGRA, *Phys. Rev. D* **109**, 044066 (2024).
- [130] S. Morisaki and V. Raymond, Rapid parameter estimation of gravitational waves from binary neutron star coalescence using focused reduced order quadrature, *Phys. Rev. D* **102**, 104020 (2020).
- [131] M. Vallisneri, J. Kanner, R. Williams, A. Weinstein, and B. Stephens, The LIGO open science center, *J. Phys. Conf. Ser.* **610**, 012021 (2015).
- [132] V. Tiwari, Estimation of the sensitive volume for gravitational-wave source populations using weighted Monte Carlo integration, *Classical Quantum Gravity* **35**, 145009 (2018).
- [133] gwosc.org.