Optimized search for a binary black hole merger population in LIGO-Virgo O3 data

Praveen Kumar and Thomas Dent

IGFAE, Campus Sur, Universidade de Santiago de Compostela, 15705 Santiago de Compostela, Spain

(Received 29 April 2024; accepted 10 July 2024; published 26 August 2024)

Maximizing the number of detections in matched filter searches for compact binary coalescence (CBC) gravitational wave signals requires a model of the source population distribution. In previous searches using the PYCBC framework, sensitivity to the population of binary black hole (BBH) mergers was improved by restricting the range of filter template mass ratios and use of a simple one-dimensional population model. However, this approach does not make use of our full knowledge of the population and cannot be extended to a full parameter space search. Here, we introduce a new ranking method, based on kernel density estimation with adaptive bandwidth, to accurately model the probability distributions of binary source parameters over a template bank, both for signals and for noise events. We demonstrate this ranking method by conducting a search over LIGO-Virgo O3 data for BBHs with unrestricted mass ratio, using a signal model derived from previous significant detected events. We achieve over 10% increase in sensitive volume for a simple power-law simulated signal population, compared to the previous BBH search. Correspondingly, with the new ranking, eight additional candidate events above an inverse false alarm rate threshold 0.5 yr are identified.

DOI: 10.1103/PhysRevD.110.043036

I. INTRODUCTION

The catalog of detected gravitational wave (GW) signals from compact binary mergers since the start of the Advanced detector era [1–3] reached little short of 100 at the end of the O3 observing run [4]. Given the relatively low signal to noise ratio (SNRs) of such signals, currently inevitable due to extremely small GW amplitudes despite advanced technologies, such signals cannot always be unambiguously distinguished from detector noise fluctuations. The presence of non-Gaussian noise artifacts (e.g., [5]) also complicates the reliable identification of signals [6].

In drawing astrophysical inferences from such catalogs, one may wish to limit the expected overall contamination by noise events in a sample, or to impose that any given event of interest is signal rather than noise with sufficiently high confidence. The first of these objectives is addressed by the calculation of false alarm rate (FAR), based on a ranking of detected event properties (the SNR and various measures of signal consistency) from least to most signal-like. Search methods and algorithms, including the choice of this ranking, may be optimized to maximize the number of events detected with FAR less than a given threshold. As detailed in [7–9], an optimal ranking is given by (any monotonic function of) the ratio of the likelihood of the measured data under the signal hypothesis to its likelihood under the noise hypothesis. In practice, we do not know either of these likelihoods exactly, so various approximations are applied. In this work, in the context of a matched-filter binary merger search, we will consider the dependence of these likelihoods on binary masses and spins, and the consequences for search optimization, in more detail.

In such a search, in contrast to the case of a fixed, known signal, we are faced with a composite hypothesis: a binary merger GW signal has many parameters with *a priori* unknown values. These can be separated into intrinsic parameters—binary component masses and spins—and extrinsic parameters, describing the relative spacetime positions and orientations of the observer and source. To optimize search sensitivity, the likelihood ratio for the signal hypothesis vs the noise hypothesis (which in the ideal case is a Gaussian function of the matched filter SNR) must be marginalized over both sets of parameters, using suitable prior distributions. These "optimal" parameter priors in fact describe an astrophysical source population distribution [10]. We then face the issue that this source distribution is partly unknown: while the distribution over binary orientation, direction from Earth, and merger time can be found simply from symmetry arguments, the distribution over component masses and spins is

only constrained by previous observations, or by theoretical/modeling expectations.

In matched-filter searches, the technical problem of unknown source parameters is addressed by searching over large number of templates (i.e., candidate signal waveforms), arranged in order to minimize the loss of SNR for any possible signal within a given parameter range (see, e.g., [11–13]). The range, i.e., boundaries, of such a template bank have first to be set, generally by considering both the prior range of possible signals, and technical constraints such as computational cost: for instance, searches for binaries with component masses below $1M_\odot$ require very large numbers of templates [14]. To optimize search sensitivity for a given population, we should then "weight" each template (adjusting its ranking at a given SNR) using its relative rate of detections, i.e., by the density of detectable signals over the bank parameters, divided by the density of templates [10]. Broadly, our aim in this work is to provide an approximate estimate of the signal density via a kernel density estimation (KDE) of previously identified, significant merger events. The combination of the signal KDE with a KDE of the search templates produced with the same method will yield the optimal weighting, as far as it can presently be determined.

The relative ranking of events occurring in different templates serves two apparently distinct goals. First, to maximize the search sensitivity, i.e., the expected number of detected signals at a fixed FAR threshold. Second, returning to the question of confidence in an individual event, we wish the ranking to indicate accurately the relative probabilities of signal (astrophysical) vs noise (terrestrial) origin: a given ranking statistic value should correspond to a constant odds of signal vs noise. Both of these goals are ideally achieved if the ranking accounts for the true distribution of signals, but as this is unknown, our strategy is to substitute a best estimate based on previous observations.

In searches performed before or soon after the first GW signal detection [15], there was little or no direct, or even indirect indication of the true binary mass and spin distribution, beyond the expectation of some nonzero rate of binary neutron star (BNS) and neutron–star–black-hole mergers [16]. Hence, either a "uninformative" or "template-agnostic" ranking was used, effectively treating each template on an equal basis, or templates were weighted only on the basis of their different responses to detector noise transients [17–19].

As pointed out in [10], such rankings without a prior weighting are implicitly optimized for a distribution of signals (above a given SNR threshold) which mirrors that of templates. Given the frequency dependence of the binary signal phase and of detector sensitivities, the distribution of templates is heavily skewed towards much higher densities at low values of chirp mass $\mathcal{M} = (m_1 m_2)^{3/5}(m_1 + m_2)^{-1/5}$, where m_1 and m_2 are the binary component masses. The density of templates over component spins is also generally somewhat higher towards extreme (near unity) dimensionless spin magnitudes. Both of these trends are very different from the population of detectable merging binaries that has emerged in more recent observations, which is rather centered on moderate or large component masses from $\sim 10M_\odot$ through $\sim 60M_\odot$ or more, with generally low (of the order of 0.1) effective orbit-aligned spins [20,21].

The template-agnostic ranking used in searches up to the second LIGO-Virgo observing run (O2) [22] thus did not optimize detection probability, either for the actual population of merger signals or an approximate population model. Various strategies for improvement were deployed to produce event catalogs from O3: a ranking incorporating a prior signal distribution uniform over $\log m_1$, $\log m_2$ was used in the GstLAL search [4,23–25], accounting for template density effects, thus maintaining broad search sensitivity, though not necessarily a realistic estimate of the astrophysical population. PYCBC-based searches pursued a different strategy: both a template-agnostic search over a broad parameter space, and a separate binary black hole (BBH) search over a restricted range of template component masses ($m > 5M_\odot$) and mass ratios ($q \equiv m_2/m_1 > 1/3$) were run, with the BBH search using a ranking that approximates the relative densities of signals and templates as a power of $\mathcal{M}$ [4,24]. While this hybrid search strategy—comparable to approaches used for independent searches of open data [26–30]—yields generally high sensitivity, it is unduly complicated: such setups imply abrupt (step-function) changes of the signal and template densities over binary mass parameters, and do not readily allow for more realistic modeling of the signal distribution.

Searches using template-agnostic rankings also generally result in biased estimates of probability of astrophysical origin $p_{\rm astro}$. An alternative approach pursued by [31] for O3 catalogs [4,25] applied a broad population model to the estimate of $p_{\rm astro}$, but without using this model in the search ranking. For the PYCBC search, such bias was partly addressed by estimating signal and noise event densities within separate template bins [4,25]. This binning approach, though, implies piecewise-constant estimates of signal density, which we do not expect to be a good fit to the actual distribution; moreover, a binned estimate over more than one dimension is likely to be dominated by counting uncertainties, given the relatively small number of confidently detected events.

Here, we pursue more sophisticated methods, which will ultimately allow a unified approach to candidate ranking within a single search. We aim to retain the advantages in sensitivity and accuracy of $p_{\rm astro}$ estimation of a statistic that incorporates a realistic population prior over stellar-mass merging binaries, without imposing arbitrary or hard boundaries within a broad search space. Ranking statistics

incorporating prior population models have already been applied in a range of other scientific contexts, including a focused search for BNS coalescences [32], search for binaries with subsolar-mass components [33], and searching for gravitationally lensed counterpart signals of known BBHs [34].

This paper is structured as follows: Section II explains how we calculate KDEs in a three-dimensional space for detected signals and template banks, and how the resulting density estimates are implemented in the PYCBC ranking statistic. Section III describes our search for BBH using this ranking in LIGO-Virgo O3 data: in Sec III A we explain the analysis configuration, in Sec. III B we compare its sensitivity to previous PYCBC searches, and in Sec. III C we give complete search results from our analysis and compare these with previous candidate lists. Section IV summarizes the implications of our results and discusses further technical issues and possible developments.

II. KDE-BASED RANKING STATISTIC FOR THE PYCBC SEARCH

A. Technical context

The PYCBC offline search pipeline, built using the PYCBC software package, [19,35–37], plays an important role in identifying candidate GWs from CBCs. It uses a matched filtering technique to correlate observed data with template waveforms to calculate the SNR. Triggers are generated in each detector separately by applying a threshold and clustering to the resulting SNR time series. To improve the distinction between CBC signals and non-Gaussian noise transients, a χ^2 test is performed for each trigger [38]. The matched-filter SNR is reweighted by the value of the χ^2 statistic; for signals matching our templates, the reduced χ^2 is expected to be near unity.

The search checks for any triggers during periods of instrumental or environmental artifacts and vetoes them. For a candidate event to be considered, it must have a consistent arrival time in different detectors, within a GW travel time window. The triggers also need to match the same template in each detector, ensuring consistency. To minimize false alarms, a coincidence test is conducted across all detectors, verifying parameter consistency [37]. Triggers passing both time and parameter coincidence tests are labeled as candidate events.

The next step of the pipeline is to measure the FAR as a function of detection statistic and use this to assign a statistical significance to candidate events. It is measured empirically by artificially shifting the time of triggers in one detector compared to another and then identifying the resulting (nonphysical) coincident events. Repeating this analysis for many different choices of time shift creates a big set of background data that is employed to estimate the FAR of the search. Since different templates in the bank can respond differently to detector noise, the search background is not the same across all templates. The measured noise background is thus fitted as a function of the template parameters, to account for the variations of the noise distribution across the target signal space [19].

The ranking statistic is a fundamental aspect of the search pipeline, offering a systematic approach to evaluating the significance of candidate coincidences. It is a function that incorporates an array of factors such as matched filter SNR, χ^2 values from signal-based vetoes, as well as intrinsic and extrinsic parameters characterizing potential GW signals [19]. It serves as a crucial tool for distinguishing astrophysical GW events from background noise. The calculation of significance also requires that candidate events are statistically independent. However, both noise and signals can create multiple triggers across the bank, leading to several related events in a short time. To ensure independence, the pipeline performs a final step of clustering: if more than one event happens within a fixed time window, only the one with the highest detection value is chosen as a candidate. The same clustering step is also applied to events in the time-shifted analyses. To maximize our ability to distinguish GW signals in the detector data, it is crucial to have a ranking statistic that reflects our knowledge of the signal distribution. As discussed in Sec. I, this ensures a reliable indication of the relative probabilities between signals and noise.

Here, we introduce a new method to assign ranking statistic based on KDE with adaptive bandwidth selection. This method allows for more accurate estimation of the rate densities of signal and noise triggers, which are then used to calculate the optimal or likelihood ratio statistic as follows:

$$\Lambda_{SN}(\vec{k}) = \frac{r_S(\vec{k})}{r_N(\vec{k})}, \tag{1}$$

where $r_S(\vec{k})$ and $r_N(\vec{k})$ are the rate densities of signal and noise triggers over $\vec{k}$, respectively, and the vector $\vec{k}$ includes an event's intrinsic and extrinsic parameters:

$$\vec{k} = \{[\rho_a, \chi^2_a, \sigma_a], \vec{\theta}, [\mathfrak{A}_{ab}, \delta t_{ab}, \delta\phi_{ab}]\}. \tag{2}$$

Here ρ_a, χ^2_a, and σ_a are the trigger SNR, signal-glitch discriminators, and template sensitivity, respectively, for each participating detector, denoted by a. Template intrinsic parameters, $\vec{\theta}$, comprise the binary component masses and spins. The network consistency parameters $\mathfrak{A}_{ab}, \delta t_{ab}$, and $\delta\phi_{ab}$ are the amplitude ratio, time difference, and phase difference, respectively, between each detector pair labeled by $a \neq b$. We write the signal rate density as $r_S(\vec{k}) \equiv \mu_S \hat{r}_S(\vec{k})$, where μ_S is an astrophysical coalescence rate per volume per time, assumed to be constant. To account for the vast dynamic range of expected rate densities, the logarithm of the ratio of signal and noise

rate densities has been used [9]; thus, we obtain a ranking statistic,

$$\mathfrak{R} = \log \hat{r}_S(\vec{k}) - \log r_N(\vec{k}), \tag{3}$$

which is equal to $\log \Lambda_{SN}$ up to a constant. The dependencies of $\mathfrak{R}$ on $\rho_a, \chi^2_a, \sigma_a$ and on the network consistency parameters $\mathfrak{A}_{ab}, \delta t_{ab}, \delta\phi_{ab}$ were already estimated in the previously derived ranking of [9], which we will denote as $\mathfrak{R}_0$.

Here, we are interested in the dependence of the statistic on the template intrinsic parameters, θ. The rate of noise triggers *per template*, which depends on both the properties of the template waveform and the data, is included as a term in $\mathfrak{R}_0$, but no explicit dependence on $\vec{\theta}$ was included. As discussed in the Introduction, this choice implies that every template is equally likely to detect a signal at any given SNR [10]. With a more general choice of ranking, including an explicit model of the distribution of signals, we have

$$\mathfrak{R} = \mathfrak{R}_0(\vec{k}) + \log \frac{d_S(\vec{\theta})}{d_T(\vec{\theta})}, \tag{4}$$

where the second term is the (log of) the ratio of signal and template densities over the space of masses and spins $\vec{\theta}$.

For producing catalog results on O3 data, the PYCBC search was run with two different configurations: broad (covering the full parameter space) and BBH (focused) [24]. The broad search is designed to detect a diverse range of signals, exploring a wide spectrum of masses and spins. On the other hand, the BBH search is focused specifically on the regions with mass ratios and component masses ranging over $1/3 \leq q \leq 1$, and $5M_\odot \leq m_1 \leq 350M_\odot$ ($m_2 \geq 5M_\odot$), respectively. These two searches had various differences in configuration, most notably in the choices of ranking statistic. In the broad search using $\mathfrak{R}_0$ as statistic, $d_S(\vec{\theta})/d_T(\vec{\theta})$ is effectively set to 1; however, in the BBH search, a nontrivial weighting is used. The distribution of templates over chirp mass has an approximate dependence $\propto \mathcal{M}^{-11/3}$, thus the weighting takes the inverse of this factor, giving a ranking statistic [27],

$$\mathfrak{R}_{\rm BBH} = \mathfrak{R}_0 + \frac{11}{3}\log\left(\frac{\mathcal{M}}{\mathcal{M}_*}\right), \tag{5}$$

where $\mathcal{M}_*$ is a constant reference mass. Hence, higher mass templates receive more weight to ensure that quiet signals at high masses are not overshadowed by more numerous noise events at lower masses. As a result, this statistic was approximately optimized for a signal distribution *uniform* over $\mathcal{M}$ (at a constant SNR); compare the ranking used in [32], optimized for a BNS population following a Gaussian distribution over $\mathcal{M}$.

This O3 configuration, with two separate searches to cover the mass space, was effective at maintaining relatively high sensitivity, but has the disadvantage of complexity; there is also a difficulty in assessing the significance of any candidate, as the two searches both produce noise events, which may be partly correlated or in common with each other. Here, we would like to both simplify the configuration by running one single search, and also maintain the advantage in sensitivity that the O3 BBH search derived from its restricted range of q and its 1-d estimate of the template distribution.

To demonstrate our new approach, for computational cost reasons, we will conduct a search over only the stellar-mass BBH space with component masses $m_1, m_2 \geq 5M_\odot$ but *without* any restriction on mass ratio; the corresponding "unrestricted BBH bank" is detailed later in Sec. III A.

B. KDE evaluation and use in ranking statistic

To develop a ranking statistic across a bank covering a wider range of mass ratios, we explore the dependence of signal/template densities over mass ratio and spin, which was neglected in the O3 BBH ranking. We now estimate signal and template densities over a three-dimensional space of $\log(\mathcal{M})$, η, and $\chi_{\rm eff}$, where η represents the symmetric mass ratio $m_1 m_2/(m_1+m_2)^2$ and $\chi_{\rm eff}$ is the effective spin defined as $(s_{1z}m_1 + s_{2z}m_2)/(m_1+m_2)$. Instead of the function of $\mathcal{M}$ used in the BBH statistic, the term $d_S(\vec{\theta})/d_T(\vec{\theta})$ is given as the ratio of KDEs calculated from signal and template points. The KDE ranking is thus given as

$$\begin{aligned}\mathfrak{R}_{\rm KDE} = \mathfrak{R}_0 &+ \log d_S(\log\mathcal{M}, \eta, \chi_{\rm eff}) \\ &- \log d_T(\log\mathcal{M}, \eta, \chi_{\rm eff}).\end{aligned} \tag{6}$$

The three-dimensional KDEs thus incorporate a more accurate model of the signal and template/noise distributions over masses and spins than for O3 BBH, with signal and template distributions modeled over chirp mass only. This not only allows us to calculate a ranking statistic more accurately, helping us better identify and understand astrophysical sources, but also serves our primary purpose of maximizing sensitivity or increasing the number of detections.

To evaluate $d_{S,T}(\log\mathcal{M}, \eta, \chi_{\rm eff})$, we use adaptive width KDE (awKDE) in preference over the fixed global bandwidth KDE, because the latter tends to overestimate or underestimate the width of features in the distribution, which may neglect structures of interest or introduce unphysical gaps in the estimated distribution. Fixed bandwidth KDE struggles to both accurately reconstruct small-scale features in densely populated regions, and supply sufficient smoothing to avoid artifacts in sparsely populated regions [39]. In contrast, awKDE addresses these limitations by dynamically adjusting the bandwidth according to

local point density, allowing for a more accurate estimation. A similar awKDE technique for the one- and two-dimensional cases was demonstrated by [39] and we extend the same approach to three dimensions. The algorithm calculates an awKDE estimator $\hat{f}$, using as input the measured parameters $\vec{X}_i$ of a set of independent events labeled by i, via [40]

$$\hat{f}(\vec{x}) = \frac{1}{n}\sum_{i=1}^{n}\frac{1}{h\lambda_i}K\left(\frac{\vec{x}-\vec{X}_i}{h\lambda_i}\right), \tag{7}$$

where K is the standard (in general multivariate) Gaussian kernel,

$$K(\vec{y}) = \frac{1}{\sqrt{2\pi}}\exp\left(-\frac{1}{2}|\vec{y}|^2\right), \tag{8}$$

n is the total number of samples, and the term $h\lambda_i$ accounts for a local bandwidth. For the adaptive estimator, we aim to take into account the local density of events when choosing the kernel bandwidth: regions with higher (lower) densities will have narrower (broader) kernels applied. This helps to avoid excess variance artifacts in regions with low event density, by applying more smoothing there, and conversely provides higher precision in regions with high event density. To implement the adaptive KDE, a pilot density $\hat{f}_0$ is first calculated by setting $\lambda_i = 1$ for all i, i.e., a standard fixed bandwidth KDE. Then based on pilot density, the local bandwidth parameter is defined as

$$\lambda_i = \left(\frac{\hat{f}_0(X_i)}{g}\right)^{-\alpha}, \qquad \log\, g = \frac{1}{n}\sum_{i=1}^{n}\log\hat{f}_0(X_i), \tag{9}$$

where α is the local bandwidth sensitivity parameter, with a value $0 < \alpha \leq 1$, and g is a normalization factor. The method requires the initial global bandwidth h and sensitivity parameter α values to be set: we assigned these using a 2-d grid search, with the figure of merit to be maximized being the total likelihood of the training samples evaluated via cross validation [39].

C. Validation of the KDE statistic

We demonstrate the effect of the statistic as a proof of principle for a coarse BBH template bank covering an unrestricted range of mass ratios. We used a template bank from the 3-OGC search [28]: that search was carried out with a bank constructed in four parts using a stochastic placement method [41]. Out of these, we use the "All BBH" bank, which is laid out with a minimal-match value of 0.97 (see Fig. 2 of [28]); however, we removed all templates with component masses $m_1, m_2 > 5M_\odot$, as the great majority of BBH events have higher masses, and the computational cost of the study is greatly reduced by this cut. The signal events used are from the 3-OGC analysis [28], covering up to the first half of the O3 run; as for the template bank, we impose the condition $m_1, m_2 > 5M_\odot$ on signal component masses, and additionally apply a threshold of IFAR > 0.5 yr to select only significant signal events, as also in Sec. II D. We thus construct KDEs using a total of 11,207 templates, or 39 signal events, respectively.

We analyzed O3 data from 2019-04-01 to 2020-01-13 in order to compare results with the O3BBH and KDE ranking statistics. Detailed search configuration settings are discussed in Sec. III A, and the suite of simulated signals (injections) analyzed to estimate sensitivity is described in detail in Sec. III B. A comparison of search sensitivity over bins in chirp mass and mass ratio, at various IFAR thresholds, is shown in Fig. 1. We observed higher sensitivity with the KDE statistic both when binned over $\mathcal{M}$ and q, particularly at higher significance thresholds, except at lower q values where the KDE statistic gives slightly lower sensitivity, within the uncertainty due to a finite set of injections. This trend is expected, as the KDE statistic is designed to down-rank parameter ranges which are less represented (or absent) in the signal training set.

D. Configuration for optimized O3 search

We now discuss specific choices in our calculation of signal and template KDEs used for the optimized search presented in the next section. For the signal case, we use events from the 4-OGC analysis of data up to the end of O3 [42], deriving our training points from the median parameter estimation (PE) source masses and applying a factor $(1+z)$ based on the PE median redshifts. For the template case, we use a stochastic BBH bank with unrestricted mass ratio; all the templates have $m_2 > 5M_\odot$ and $m_1 > m_2$ (again corresponding to detector frame, i.e., redshifted, masses), extending up to masses of $500M_\odot$; this bank is described in detail in Sec. III A. In the signal case, an additional condition is applied, requiring IFAR > 0.5 yr: we then impose the same minimum component mass of $5M_\odot$, resulting in a total of 59 events. These conditions ensure that the events used in the analysis are BBH-like. We then evaluated the KDE using (7) over the three-dimensional space of $\log(\mathcal{M})$, η, and $\chi_{\rm eff}$ parameters of the template bank for both signal and template cases.

We show the resulting KDEs in Fig. 2 for both the signal (top) and template (bottom) cases across different 2-d parameter space slices. In the signal KDE (top), there is a higher density of events concentrated around $\eta \approx 0.25$, $\chi_{\rm eff} \approx 0$, and $\mathcal{M} \approx 40M_\odot$, while for the template KDE (bottom) we observe various different trends; for instance, the density of templates is higher towards nonzero $\chi_{\rm eff}$ values (both positive and negative). Scatter plots of the ratio of signal to template KDEs at the template points are shown in Fig. 3: These reflect that the highest density of detected events is at high values of $\mathcal{M}$ and η and near $\chi_{\rm eff} = 0$.

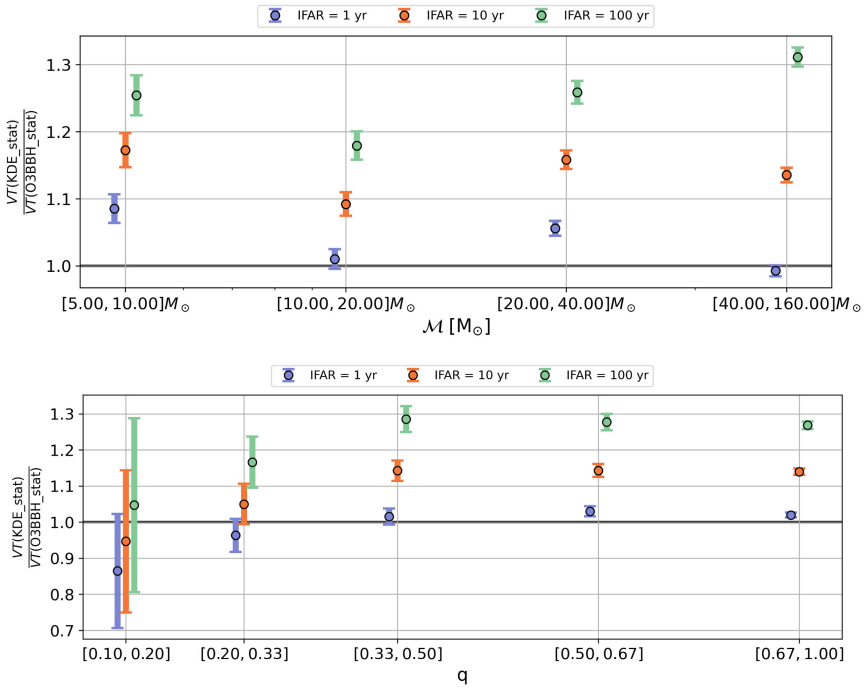

FIG. 1. Direct comparison of search sensitivity between two ranking statistics: KDE and O3 BBH, for the same bank and the same set of injections. The top and bottom subplots show $\langle VT \rangle$ comparisons for injections divided into chirp mass bins and mass ratio bins, respectively, at different inverse false alarm rate (IFAR) thresholds.

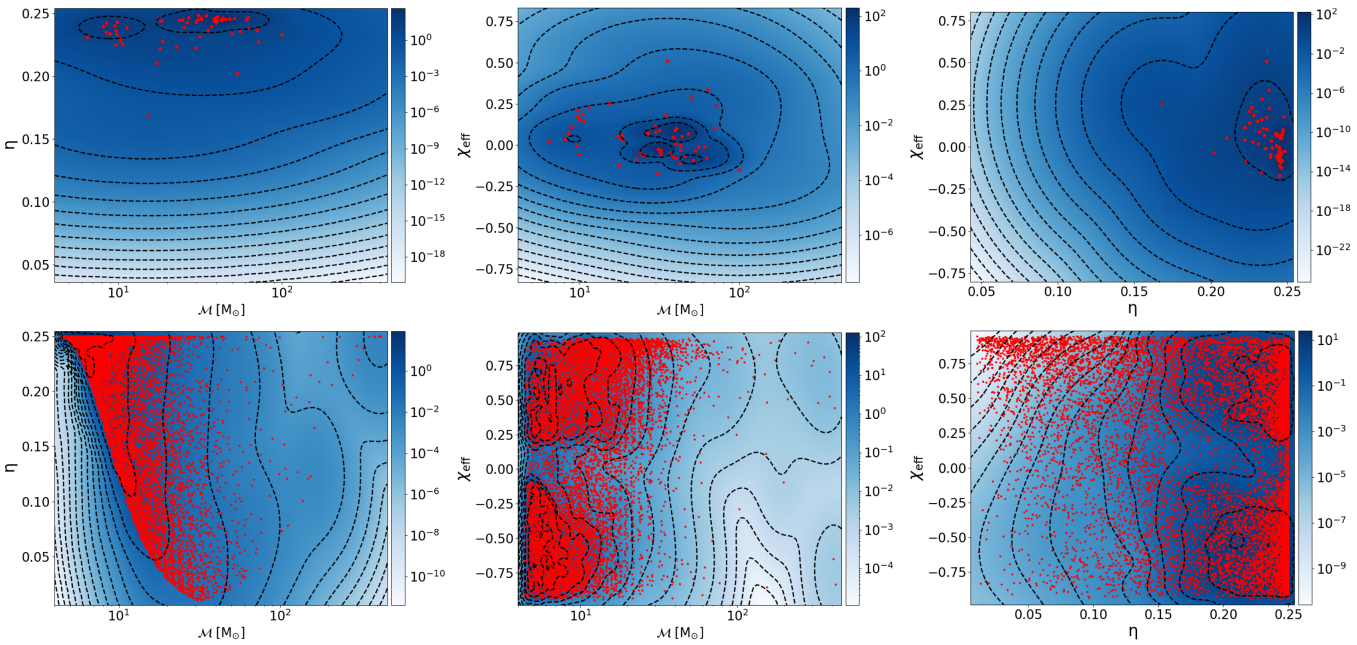

FIG. 2. Top/bottom row: The signal, respectively template KDE, plotted over various 2-d "slices" through our 3-d parameter space. The red dots represent signals and templates in the signal and template KDE, respectively, used for training. The values of the third parameter for these slice plots are $\chi_{\rm eff} = 0$, $\eta = 0.25$, and $\log(\mathcal{M}) = 2$ in both rows, respectively.

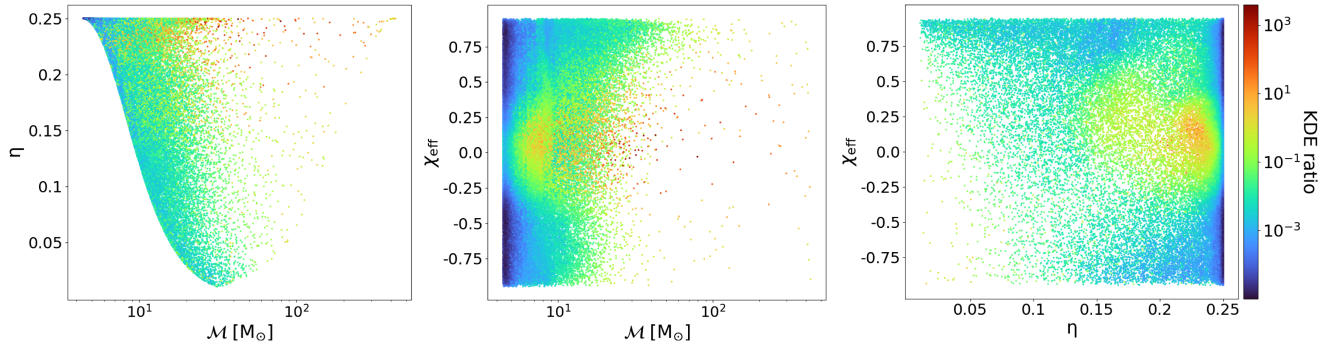

FIG. 3. Scatter plots between the different parameters used in the KDE. The color bar is a ratio of signal to template KDE, evaluated at the template points.

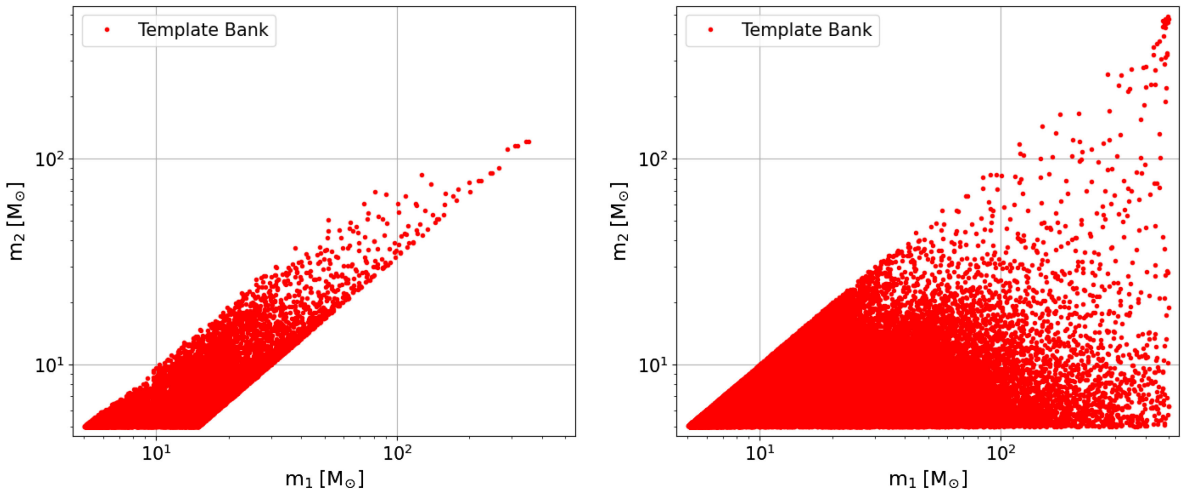

FIG. 4. Template placement for the O3 focused BBH bank with restricted q (17,094 templates, left) and our BBH bank with unrestricted q (64,184 templates, right).

III. SEARCH OF THE LIGO-VIRGO O3 DATA

A. Search configuration and sensitivity

To maintain high sensitivity while conducting only one search for BBH-like systems, in contrast to the previous PYCBC search strategy in O3, we employed an unrestricted BBH bank in this work, covering a broad range of component masses (5 to 500$M_\odot$) and orbit-aligned component spins (−0.949 to 0.949). The O3 BBH search achieved high sensitivity with a small number of templates, and so a low noise background, but at the cost of leaving the small-q region to a separate search with worse template coverage, and of having two (mostly) independent sets of noise events over the two searches. Our approach aims to address these challenges by covering the entire range of q in a single search, while maintaining sensitivity to the near-equal-mass BBH population.

Our bank is generated using a brute force stochastic method with a minimal-match value of 0.985 [41], employing the SEOBNRv4_ROM waveform approximant for templates, and using the same power spectral density estimate as in the 3-OGC search of O3a [28]. The chosen minimal match strikes a good balance between optimizing the computational speed and minimizing SNR loss. While it might seem logical to increase the value to reduce SNR loss risk, the major drawback of raising it is the increased computational cost: as the value goes up, the number of templates in the bank also increases. Thus, a balance is needed between the template bank's quality and the computational resources required. Figure 4 (right) shows our unrestricted BBH bank in the $m_1 - m_2$ plane, for comparison with the O3 focused BBH bank (left).

We analyze the complete set of data from the O3 run, recorded and calibrated by the LVK collaborations, which is also accessible via the Gravitational Wave Open Science Center (GWOSC)[1]: we used offline recalibrated (C01) strain data, applying available vetoes for invalid or sub-standard data quality [5]. Given the new unrestricted bank, we used the configuration settings appropriate to the BBH

[1]https://gwosc.org/GWTC-3/

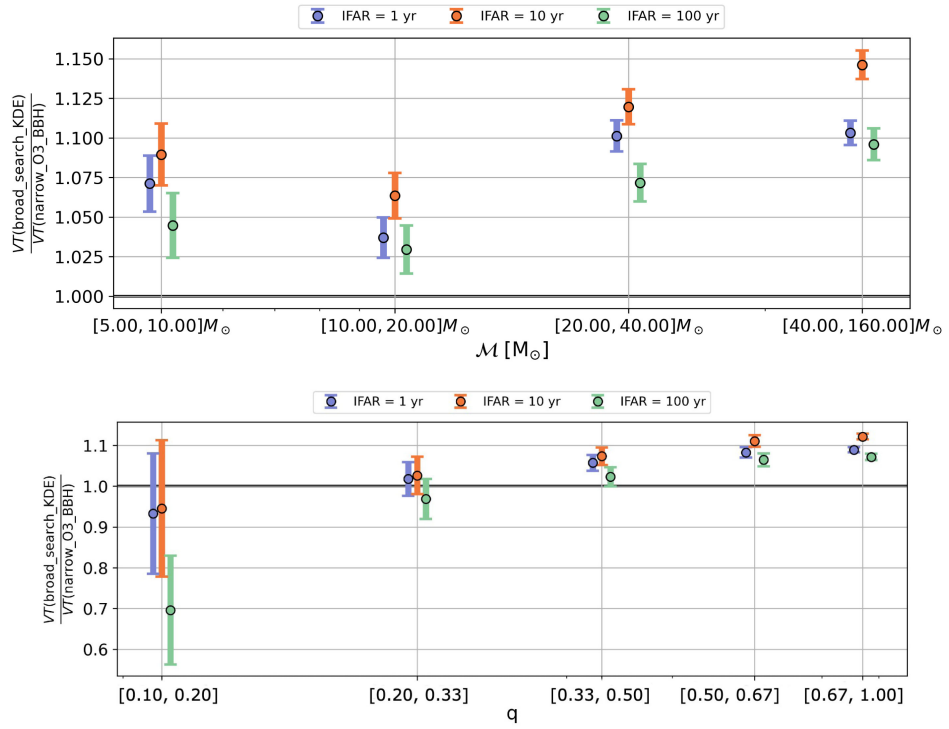

FIG. 5. Comparison of search sensitivity between two searches: KDE (broad) and O3 BBH (narrow), using the same set of injections. The top and bottom subplots show $\langle VT \rangle$ comparisons for injections divided into chirp mass bins and mass ratio bins, respectively, at different IFAR thresholds.

search from [4], with slight refinements. We used a slightly different χ^2 tuning which makes sure that no template uses less than 11 χ^2 bins [42], and a low-frequency cutoff f_{low} of 10 Hz is applied for the matched filter SNR evaluation across all detectors. The observation of repeated (lower amplitude) "echo" glitches, a few seconds away from loud autogated glitches which are windowed to zero [37], prompted the implementation of an additional veto [43]. This veto is applied to all triggers between 1.5 seconds before and 2.5 seconds after the central times of the gates for H1 and L1. For V1, triggers within 0.5 seconds before and 1 second after the central gate times are excluded. These time intervals are selected based on visual inspection of the distribution of single-detector triggers around gated times. Following these settings, the PYCBC search is run with the KDE ranking statistic.

B. Search sensitivity comparisons

To estimate the sensitive volume-time product $\langle VT \rangle$ of the searches, a large set of injections (simulated signals) are added to the data; we then count how many are recovered, i.e., detected by the search with an IFAR above a specified threshold. The anticipated number of detections for a search is $\hat{N} = \langle VT \rangle R$, where R denotes the astrophysical rate of mergers per unit volume and unit observing time. The injection component masses and spins are distributed within the range of 2 to $100M_\odot$, and between -0.998 to 0.998, respectively; their distribution, detailed in [4,44], is primarily characterized by favoring more equal masses. The mass distribution follows separate power laws for each component mass: $p(m_1) \sim m_1^{-2.35}$ and $p(m_2|m_1) \sim m_2$, while the redshift distribution follows $p(z) \sim dV_c/dz$, where $V_c(z)$ is the comoving volume out to redshift z.[2] Injections are generated using the SEOBNRv4PHM waveform model [46]. Given the mass range of our bank, we consider only injections with $\mathcal{M} \geq 5M_\odot$.

Figure 5 shows a comparison of injection sensitivities between the KDE-based (broad) and O3 BBH (narrow) searches over the entire O3 dataset. These searches exhibit different characteristics, due to the nature of their respective template banks and the ranking statistic employed. In the case of O3 BBH, the templates are narrowly concentrated around equal masses [47] with mass ratios q ranging from $1/3$ to 1; in contrast, the KDE search encompasses a broader parameter space with a wide range of mass ratios. The purpose of using a broader bank is to achieve a high match, and so a high recovered SNR, for signals with a wider range of parameters. While the broader mass ratio range contributes to this sensitivity improvement, it might also result in a higher noise background, due to the larger number of templates. However, the KDE statistic

[2]This redshift distribution corresponds to a comoving rate density that *increases* as $(1+z)^1$: this factor cancels in $p(z)$ against a $(1+z)^{-1}$ factor due to time dilation [45].

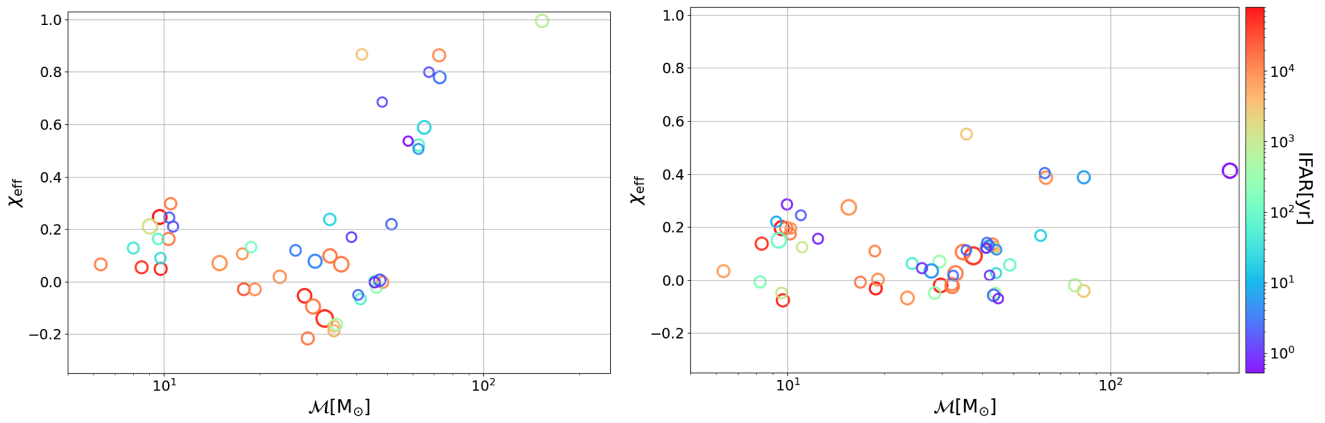

FIG. 6. Events detected with a threshold of IFAR > 0.5 yr by two different searches: O3 BBH (left) and our unrestricted BBH search using the KDE statistic (right). The color bar represents IFAR, and the size of the symbols corresponds to the network SNR.

effectively penalizes templates in regions with a high density of noise events allowing a focus on more promising areas. This enhances our ability to detect GWs, even when injections do not precisely match any astrophysical distributions. Despite injections exhibiting a relatively smooth distribution, rather than following a detailed astrophysical population model, the KDE statistic enhances sensitivity to these injections.

Figure 5 (bottom) compares injection sensitivities across mass ratio bins at various IFAR thresholds. This comparison addresses a potential question regarding whether our higher sensitivity, observed over $\mathcal{M}$ bins, can only be attributed to the bank covering a broader range of mass ratios. We observe lower sensitivity at low q values; this may be expected, as our KDE signal model disfavors unequal-mass systems. The corresponding injections may still be detected in the O3 BBH search, with biased template parameters, while they can be detected in our search with low q (or high spin) templates, resulting in a lower signal KDE and thus a lower ranking. Conversely, in the range of mass ratios that is fully covered by the O3 BBH bank, the KDE statistic still yields higher sensitivity; this must be due to the more detailed modeling of the template and signal densities by KDEs.

C. Search results

We now discuss the set of events detected by our new KDE search and compare it with other searches for BBH mergers over the O3 dataset. Initially, the PYCBC O3 BBH search used in GWTC-3 identified a total of 49 events with an IFAR exceeding 0.5 years, with a total probability of astrophysical origin ($p_{\rm astro}$) over these events of 48.33. Using the unrestricted bank and our new KDE statistic, we see a number of detections increased by over 10%, with 57 events meeting the same threshold of IFAR > 0.5 yr, and with a total astrophysical probability of 55.04; Fig. 6 illustrates the events detected by both searches. A total of 62 events have $p_{\rm astro}$ greater than 0.5 in our search. Additionally, the KDE search outperforms the 4-OGC search [42] which reported 50 BBH events with IFAR > 0.5 yr. Our event set is reported in Tables I (high purity) and II (marginal), which also includes some candidates with IFAR < 0.5 yr but having a high probability of astrophysical origin.

Since our signal model is trained on a sample of previous BBH detections, many from O3, one might suspect the result for O3 data to be biased. Specifically, the statistic might be up-ranking the exact templates which correspond to known O3 events, rather than reflecting the underlying

TABLE I. GW candidate events with either IFAR or $p_{\rm astro}$ value greater than 0.5 yr from our KDE-based broad BBH search of O3 data, sorted by observation time. The table presents a low contamination "gold sample" selected such that the cumulative $p_{\rm terr}$ is below 1 (see main text). Events not found with IFAR > 0.5 yr in previous PYCBC-based searches [4,25,28,42] are marked in bold. Template component and chirp masses are provided in detector frame.

Event	GPS	$m_1(M_\odot)$	$m_2(M_\odot)$	$\mathcal{M}(M_\odot)$	$\chi_{\rm eff}$	$p_{\rm astro}$	IFAR (yr)	SNR
190408_181802	1238782700.29	30.4	23.9	23.4	−0.05	1.000	8.3×10^3	14.09
190412_053044	1239082262.17	34.9	9.60	15.3	0.27	1.000	8.3×10^3	18.26
190413_052954	1239168612.50	51.4	46.6	42.6	0.12	0.994	2.5	8.66
190413_134308	1239198206.74	89.5	61.5	64.4	0.41	0.996	2.7	9.16
190421_213856	1239917954.26	67.7	37.7	43.6	−0.05	1.000	4.8×10^2	10.12

(Table continued)

TABLE I. *(Continued)*

Event	GPS	$m_1(M_\odot)$	$m_2(M_\odot)$	$\mathcal{M}(M_\odot)$	$\chi_{\rm eff}$	$p_{\rm astro}$	IFAR (yr)	SNR
190503_185404	1240944862.30	51.4	46.6	42.6	0.12	1.000	53.0	11.83
190512_180714	1241719652.42	24.5	18.9	18.7	−0.00	1.000	9.1×10^3	12.17
190513_205428	1241816086.75	34.8	31.8	28.9	−0.04	1.000	3.2×10^2	11.72
190517_055101	1242107479.83	48.0	36.4	36.3	0.57	1.000	3.0×10^3	10.36
190519_153544	1242315362.40	89.5	61.5	64.4	0.41	1.000	9.1×10^3	13.40
190521_030229	1242442967.46	101.	82.6	79.4	−0.03	1.000	6.3×10^2	12.94
190521_074359	1242459857.47	57.4	33.7	38.0	0.10	1.000	4.3×10^4	24.36
190527_092055	1242984073.79	42.3	34.2	33.1	0.02	0.937	2.6	8.37
190602_175927	1243533585.10	101.	82.6	79.4	−0.03	1.000	2.7×10^3	12.33
190630_185205	1245955943.18	48.4	21.4	27.5	0.02	1.000	5.4	15.44
190701_203306	1246048404.58	67.7	37.7	43.6	−0.05	1.000	2.6	11.53
190706_222641	1246487219.33	115.	82.0	84.2	0.37	1.000	6.3	12.74
190707_093326	1246527224.17	12.5	10.3	9.89	−0.06	1.000	5.3×10^4	12.93
190719_215514	1247608532.93	45.5	37.0	35.7	0.12	0.937	2.3	8.09
190720_000836	1247616534.71	14.6	9.95	10.4	0.19	1.000	1.3×10^4	10.64
190725_174728	1248112066.47	21.5	5.43	9.00	0.21	0.977	9.1	9.42
190727_060333	1248242631.99	51.4	46.6	42.6	0.12	1.000	1.3×10^4	11.52
190728_064510	1248331528.53	14.4	9.60	10.2	0.20	1.000	1.3×10^4	13.24
190731_140936	1248617394.64	51.4	46.6	42.6	0.12	0.937	2.8	7.86
190803_022701	1248834439.88	64.0	38.6	43.0	0.02	0.992	27.0	8.69
190828_063405	1251009263.76	43.7	30.7	31.8	−0.02	1.000	1.4×10^4	15.90
190828_065509	1251010527.89	24.0	16.7	17.4	−0.02	1.000	1.4×10^4	10.53
190915_235702	1252627040.70	43.7	30.7	31.8	−0.02	1.000	1.2×10^4	13.07
190924_021846	1253326744.84	10.0	5.58	6.46	0.05	1.000	6.6×10^3	12.38
190925_232845	1253489343.13	22.4	20.7	18.8	0.12	1.000	1.1×10^4	10.06
190930_133541	1253885759.24	15.7	8.30	9.85	0.18	1.000	8.2×10^3	10.11
191105_143521	1256999739.93	12.7	9.59	9.60	−0.03	1.000	1.2×10^3	9.94
191109_010717	1257296855.23	34.0	32.1	28.8	0.05	0.999	3.3×10^2	12.41
191126_115259	1258804397.63	18.5	9.11	11.2	0.24	0.924	2.0	8.54
191127_050227	1258866165.56	51.4	46.6	42.6	0.12	0.973	5.2	8.25
191129_134029	1259070047.20	13.6	7.22	8.53	0.13	1.000	4.1×10^4	12.88
191204_110529	1259492747.54	41.5	23.0	26.7	0.06	0.816	0.97	8.97
191204_171526	1259514944.09	14.4	8.75	9.71	0.19	1.000	8.1×10^4	17.20
191215_223052	1260484270.34	35.0	23.9	25.1	0.08	1.000	53.0	10.42
191216_213338	1260567236.48	10.4	10.4	9.04	0.13	1.000	1.2×10^2	18.33
191222_033537	1261020955.12	92.7	38.0	50.6	0.05	1.000	1.0×10^2	11.55
191224_043228	1261197166.15	20.8	9.71	12.2	0.15	0.860	0.83	8.84
191230_180458	1261764316.41	90.5	57.0	62.2	0.16	0.999	19.0	10.09
200128_022011	1264213229.91	51.4	46.6	42.6	0.12	1.000	9.9×10^3	9.23
200129_065458	1264316116.42	34.8	31.8	28.9	−0.04	1.000	5.9×10^4	16.08
200202_154313	1264693411.56	10.2	8.51	8.12	−0.01	1.000	1.7×10^2	11.05
200208_130117	1265202095.95	51.4	46.6	42.6	0.12	1.000	1.2×10^3	9.75
200219_094415	1266140673.20	51.4	46.6	42.6	0.12	1.000	3.3×10^3	9.54
200220_124850	1266238148.16	64.0	38.6	43.0	0.02	0.846	1.1	7.96
200224_222234	1266618172.40	45.5	37.0	35.7	0.12	1.000	1.3×10^4	18.09
200225_060421	1266645879.40	21.1	20.6	18.2	−0.05	1.000	7.2×10^4	12.51
200311_115853	1267963151.39	42.3	34.2	33.1	0.02	1.000	1.3×10^4	17.74
200316_215756	1268431094.16	17.3	8.81	10.6	0.13	1.000	1.1×10^3	9.50

TABLE II. "Silver" sample of marginal BBH candidate events from O3 data, selected as in Table I but with cumulative $p_{\rm terr} \geq 1$.

Event	GPS	$m_1(M_\odot)$	$m_2(M_\odot)$	$\mathcal{M}(M_\odot)$	$\chi_{\rm eff}$	$p_{\rm astro}$	IFAR (yr)	SNR
190514_065416	1241852074.85	67.7	37.7	43.6	−0.05	0.583	0.21	8.07
190711_030756	1246849694.65	67.7	37.7	43.6	−0.05	0.558	0.25	8.91
190916_200658	1252699636.90	67.7	37.7	43.6	−0.05	0.804	0.69	7.95
190926_050336	1253509434.08	51.4	46.6	42.6	0.12	0.804	0.70	7.62
190929_012149	1253755327.51	109.0	50.7	63.9	0.09	0.669	0.39	9.20
191103_012549	1256779567.53	12.8	10.6	10.1	0.30	0.803	0.76	9.25
191208_080334	1259827432.85	45.5	37.0	35.7	0.12	0.533	0.14	7.59
191225_215715	1261346253.87	337.0	213.0	232.0	0.39	0.440	0.52	17.66
200214_223306	1265754805.00	90.5	57.0	62.2	0.16	0.647	0.32	7.80
200301_211019	1267132237.66	29.4	18.1	19.9	−0.17	0.666	0.38	8.29

population distribution. However, examining the signal KDE in Fig. 2, it varies smoothly and slowly over the region where many detections are present; thus, it is not "overfitting" to individual O3 events, and we expect the ranking to be stable under small changes in the signal training set. The fact that the search is able to detect several events that are *not* in the training set, coupled with the increased sensitivity for an independent injection set, confirms a real sensitivity improvement.

We find a conspicuous concentration of events, as well as higher sensitivity (see Fig. 5) near $\mathcal{M} \approx 40M_\odot$, consistent with the higher event count being a consequence of improved sensitivity. Several events previously reported with IFAR values below 0.5 yr now surpass this threshold: 190527_092055, 190725_174728, 190916_200658, 190926_050336, 191127_050227, 191224_043228, and 200220_124850. However one BBH event (200209_085452) reported in [21] is not observed in our results with IFAR > 0.5 yr; we find a lower significance, possibly due to elevated χ^2 values indicating a relatively poor match to available templates. Events 190514_065416, 191208_080334, and 200301_211019 from GWTC-3 also remain with IFAR below 0.5 yr.

Now considering events that were not seen in previous PYCBC BBH searches, 190711_030756 and 200214_223306 were first reported in [48] and [49], respectively. The IMBH marginal candidate event 191225_215715, first reported in [50], is also identified in our search, although with a low (<0.5) probability of astrophysical signal origin. It stands out from others as its masses are significantly above the BBH-like events we use for training the signal KDE. The signal KDE is still nonzero outside the mass range of previously detected BBHs, but the exact values that determine the ranking of this event are strongly dependent on the bandwidth and other configuration choices. Since there is no real constraint on the signal model beyond the trained region, the significance estimate for 191225_215715 may be affected by a—currently unknown—systematic bias. Adjustments to the KDE statistic to allow for possible detections outside the previously observed population range are discussed further in Sec. IV.

A nonzero $\chi_{\rm eff}$ indicates the definite presence of one or more component spins in the system, either aligned with the orbital angular momentum if positive, or antialigned if negative. Most detected binaries up to O3 are consistent with $\chi_{\rm eff} \approx 0$ [4,25]. However, several events are identified in the O3 BBH search by templates with high (positive) $\chi_{\rm eff}$, which appears inconsistent with PE analyses, or at least indicates a significant bias in the search template parameters. Using the KDE ranking statistic which incorporates an estimate of the signal spin distribution, we find that events now tend to cluster at lower $\chi_{\rm eff}$ values, narrowing the overall range of values. Thus, in addition to optimizing search sensitivity, the KDE-based ranking may also improve the accuracy of the search template masses and spins in comparison to the credible regions found by PE analysis.

1. Probability of astrophysical origin and catalog purity

Given the list of candidate GW events from the search, together with the estimated distributions of noise events and astrophysical signal events over the search ranking statistic, we apply Bayesian inference on a mixture model, with the total number of signals present in the search results (including possible high-FAR events) as an unknown parameter [51–53]. Using this signal rate, and marginalizing over its uncertainty, we calculate the probability of astrophysical origin $p_{\rm astro}$ and of terrestrial (noise) origin $p_{\rm terr} = 1 - p_{\rm astro}$ for each candidate. Highly significant events with low FAR thus have $p_{\rm terr} \ll 1$. We calculated the cumulative $p_{\rm terr}$, i.e., the cumulative sum of $p_{\rm terr}$ values, sorted in *increasing* order, over all our detected events. This cumulative sum provides an overall idea of the likely noise contamination in a given set of events when looked at together, in contrast to the individual $p_{\rm terr}$ values. These cumulative $p_{\rm terr}$ values can be used to define event sets with more or less high purity, defined as the probable fraction of events that are astrophysical signals.

Thus, given the values shown in Fig. 7, we defined two event sets based on significance and overall purity. In both tables, events are selected based on the criteria of IFAR exceeding 0.5 years or a $p_{\rm astro}$ value greater than 0.5. Table I

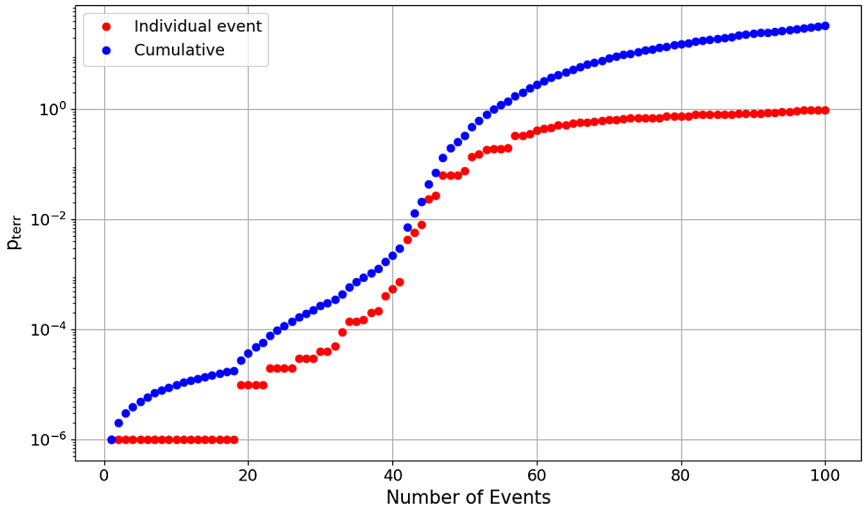

FIG. 7. Probability of terrestrial (noise) origin for the top 100 GW event candidates. The red curve indicates the individual p_{terr} values, sorted in increasing order; lower values are considered more significant. The blue curve shows the cumulative terrestrial probability up to a given event, indicating the overall purity of the set of events.

then contains events with the smallest p_{terr} values, up to a cumulative $p_{\mathrm{terr}} < 1$, which we may call a "Gold" event set. Table II then includes the remaining events for which the cumulative $p_{\mathrm{terr}} \geq 1$, known as the "Silver" event set. This classification provides valuable insights into the observations over O3, helping us define event sets of well-defined purity which may be used for further investigation into astrophysical or cosmological properties.

IV. DISCUSSION

In this paper we demonstrated a ranking statistic for gravitational-wave searches targeting compact binary mergers, which builds on previous work by using a model of the signal distribution over binary masses and spins based on previously detected significant events. The primary objective of this work is to combine the previous "broad" and "BBH focused" PYCBC searches employed for O3 data into one search only, which aims to increase sensitivity to the astrophysical population of signals, allowing for both the known high BBH signal density and the detection of new types of signals. This objective is achieved by incorporating KDE-based models of the signal and template distributions into the search ranking statistic. Conducting the PYCBC search using our developed model on the O3 data, we found around 10% more events and also higher sensitivity compared to the O3 BBH search.

The basic approach of Eq. (4), i.e., introducing a signal density model in the ranking, may be extended or modified in various ways. One possible extension addresses the issue that previous detections, being a finite sample of events, do not give a complete picture of the full signal distribution. Specifically, our choice of statistic will strongly down-rank any signals with parameters significantly outside the range of masses and spins covered by O1-O3 BBH detections. A possible remedy is to limit how far any template may be down-ranked, which we can achieve by adding an extra term to the signal density in order to represent possible events with parameters outside the previously seen distribution. Formally, the signal model is then a mixture of the KDE based on some set of known events with a broad, (near)constant density covering the entire search space (or bank). We may write

$$\hat{h}(x) = (1-a)\hat{f}(x) + aC, \tag{10}$$

where a is the mixture fraction of the broad component and C is a normalization constant: $C = 1/v$, where v is the volume of the parameter space covered by the template bank in the coordinates used for the KDE. The parameter a may be adjusted to give more or less weight to the detection of any "exceptional" events outside the known population.

One may also take a signal density based on a more or less realistic formal or astrophysical model, rather than empirically on previous detections. The signal model may be an arbitrary function, but in some cases can be obtained by applying a KDE: the estimate will use as input a set of *simulated* detectable events (i.e., those above a fixed SNR threshold) whose distribution follows the desired model. The simulated events may derive from an astrophysical calculation, for instance stellar population synthesis (e.g., [54]) or modeling gravitationally lensed counterparts [34], or from other sampling procedures. The KDE method is particularly applicable for complicated source distributions which cannot be written in closed form, or where only a relatively small number of sample events may be available.

The application of our optimized adaptive KDE to relatively large datasets (order 10^5 points or more) can be computationally demanding. Determining the best choice of bandwidth and adaptive parameter by grid search

requires some tens of KDE evaluations, becoming impractical if the time for a single KDE on a CPU core exceeds about an hour. While parallelization could be of benefit, some change in method is required for the template banks used in current searches over a full parameter space of stellar mass binaries, with up to $\mathcal{O}(10^6)$ points. One possible strategy to address computational cost for large banks is downsampling the input points, with the probability of rejection being a known function of template parameters (for instance, a power of chirp mass): after the KDE is computed from a reduced set of points, the density may then be corrected for the rejection step. Alternatively, the bank may be split into a number of sub-banks with a separate KDE evaluated over each one; while this approach is easily parallelizable, care is required in combining the sub-bank KDEs in order to avoid edge effects or other artifacts.

We may also consider applying a KDE-based statistic in a context beyond the state of the art of CBC searches over binary masses and orbit-aligned spins, with templates representing only the dominant emission multipole. When the search is extended to templates including subdominant multipole emission ("higher modes"), it is necessary to construct the bank over a space with additional dimensions representing the binary orientation, in particular the orbital inclination angle [55,56]. Similarly, a search including the orbital precession effects of nonaligned spins requires a larger number of degrees of freedom in bank construction; for instance, the spin magnitude and opening angle used in [57], or the amplitudes and phases of harmonic waveform components in [58]. In both the higher modes and precession cases, the distribution of templates and of signals over the higher-dimensional space may show nontrivial structure that is difficult to predict or model analytically; the adaptive KDE then offers a possible route to optimize the sensitivity of these more complex searches.

ACKNOWLEDGMENTS

We are grateful to Florian Aubin for carefully reading an earlier version of this paper, and to the PYCBC search development team for discussions during this project, and for making details of previous O3 searches available. This work has received financial support from Xunta de Galicia (CIGUS Network of research centers), by European Union ERDF and by the "María de Maeztu" Units of Excellence Program No. CEX2020-001035-M and the Spanish Research State Agency. T. D. and P. K. are supported by Research Grant No. PID2020–118635 GB-I00 from the Spanish Ministerio de Ciencia e Innovación. This research project was made possible through the access granted by the Galician Supercomputing Center (CESGA) to its supercomputing infrastructure. The authors are also 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.

[1] J. Aasi *et al.*, Advanced LIGO, Classical Quantum Gravity **32**, 074001 (2015).
[2] F. Acernese *et al.*, Advanced Virgo: A second-generation interferometric gravitational wave detector, Classical Quantum Gravity **32**, 024001 (2015).
[3] Y. Aso, Y. Michimura, K. Somiya, M. Ando, O. Miyakawa, T. Sekiguchi, D. Tatsumi, and H. Yamamoto, Interferometer design of the KAGRA gravitational wave detector, Phys. Rev. D **88**, 043007 (2013).
[4] 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).
[5] D. Davis *et al.*, LIGO detector characterization in the second and third observing runs, Classical Quantum Gravity **38**, 135014 (2021).
[6] D. Davis, M. Trevor, S. Mozzon, and L. K. Nuttall, Incorporating information from LIGO data quality streams into the PYCBC search for gravitational waves, Phys. Rev. D **106**, 102006 (2022).
[7] R. Biswas *et al.*, Likelihood-ratio ranking of gravitational-wave candidates in a non-Gaussian background, Phys. Rev. D **85**, 122008 (2012).
[8] K. Cannon, C. Hanna, and J. Peoples, Likelihood-ratio ranking statistic for compact binary coalescence candidates with rate estimation, arXiv:1504.04632.
[9] 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).
[10] T. Dent and J. Veitch, Optimizing gravitational-wave searches for a population of coalescing binaries: Intrinsic parameters, Phys. Rev. D **89**, 062002 (2014).
[11] T. Dal Canton and I. W. Harry, Designing a template bank to observe compact binary coalescences in advanced LIGO's second observing run, arXiv:1705.01845.
[12] S. Roy, A. S. Sengupta, and N. Thakor, Hybrid geometric-random template-placement algorithm for gravitational wave searches from compact binary coalescences, Phys. Rev. D **95**, 104045 (2017).

[13] 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).
[14] B. P. Abbott *et al.*, Search for subsolar-mass ultracompact binaries in Advanced LIGO's first observing run, Phys. Rev. Lett. **121**, 231103 (2018).
[15] B. P. Abbott *et al.*, Observation of gravitational waves from a binary black hole merger, Phys. Rev. Lett. **116**, 061102 (2016).
[16] J. Abadie *et al.*, Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors, Classical Quantum Gravity **27**, 173001 (2010).
[17] 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).
[18] B. P. Abbott *et al.*, Binary black hole mergers in the first Advanced LIGO observing run, Phys. Rev. X **6**, 041015 (2016); **8**, 039903(E) (2018).
[19] 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).
[20] R. Abbott *et al.*, Population properties of compact objects from the second LIGO-Virgo gravitational-wave transient catalog, Astrophys. J. Lett. **913**, L7 (2021).
[21] R. Abbott *et al.*, Population of merging compact binaries inferred using gravitational waves through GWTC-3, Phys. Rev. X **13**, 011048 (2023).
[22] 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).
[23] H. K. Y. Fong, From simulations to signals: Analyzing gravitational waves from compact binary coalescences, Ph.D. thesis, Toronto University, 2018.
[24] 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).
[25] 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).
[26] A. H. Nitz, C. Capano, A. B. Nielsen, S. Reyes, R. White, D. A. Brown, and B. Krishnan, 1-OGC: The first open gravitational-wave catalog of binary mergers from analysis of public Advanced LIGO data, Astrophys. J. **872**, 195 (2019).
[27] A. H. Nitz, T. Dent, G. S. Davies, S. Kumar, C. D. Capano, I. Harry, S. Mozzon, L. Nuttall, A. Lundgren, and M. Tápai, 2-OGC: Open gravitational-wave catalog of binary mergers from analysis of public Advanced LIGO and Virgo data, Astrophys. J. **891**, 123 (2020).
[28] 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).
[29] 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).
[30] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, New binary black hole mergers in the second observing run of Advanced LIGO and Advanced Virgo, Phys. Rev. D **101**, 083030 (2020).
[31] N. Andres *et al.*, Assessing the compact-binary merger candidates reported by the MBTA pipeline in the LIGO–Virgo O3 run: Probability of astrophysical origin, classification, and associated uncertainties, Classical Quantum Gravity **39**, 055002 (2022).
[32] R. Magee *et al.*, Sub-threshold binary neutron star search in Advanced LIGO's first observing run, Astrophys. J. Lett. **878**, L17 (2019).
[33] K. S. Phukon, *et al.*, The hunt for sub-solar primordial black holes in low mass ratio binaries is open, arXiv: 2105.11449.
[34] A. K. Y. Li, J. C. L. Chan, H. Fong, A. H. Y. Chong, A. J. Weinstein, and J. M. Ezquiaga, TESLA-X: An effective method to search for sub-threshold lensed gravitational waves with a targeted population model, arXiv:2311.06416.
[35] 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).
[36] 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).
[37] S. A. Usman, A. H. Nitz, I. W. Harry, C. M. Biwer, D. A. Brown, M. Cabero, C. D. Capano, T. Dal Canton, T. Dent, S. Fairhurst *et al.*, The PYCBC search for gravitational waves from compact binary coalescence, Classical Quantum Gravity **33**, 215004 (2016).
[38] B. Allen, χ^2 time-frequency discriminator for gravitational wave detection, Phys. Rev. D **71**, 062001 (2005).
[39] J. Sadiq, T. Dent, and D. Wysocki, Flexible and fast estimation of binary merger population distributions with an adaptive kernel density estimator, Phys. Rev. D **105**, 123014 (2022).
[40] D. W. Scott, *Multivariate Density Estimation: Theory, Practice, and Visualization* (John Wiley & Sons, New York, 2015).
[41] I. W. Harry, B. Allen, and B. Sathyaprakash, Stochastic template placement algorithm for gravitational wave data analysis, Phys. Rev. D **80**, 104014 (2009).
[42] 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).
[43] K. Chandra, V. Villa-Ortega, T. Dent, C. McIsaac, A. Pai, I. Harry, G. C. Davies, and K. Soni, Optimized PYCBC search for gravitational waves from intermediate-mass black hole mergers, Phys. Rev. D **104**, 042004 (2021).
[44] LIGO Scientific, Virgo and KAGRA Collaborations, GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run—O3 search sensitivity estimates, 10.5281/zenodo.7890437.

[45] B. P. Abbott *et al.*, The rate of binary black hole mergers inferred from Advanced LIGO observations surrounding GW150914, Astrophys. J. Lett. **833**, L1 (2016).
[46] S. Ossokine *et al.*, Multipolar effective-one-body waveforms for precessing binary black holes: Construction and validation, Phys. Rev. D **102**, 044055 (2020).
[47] 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).
[48] 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).
[49] 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.
[50] R. Abbott *et al.*, Search for intermediate-mass black hole binaries in the third observing run of Advanced LIGO and Advanced Virgo, Astron. Astrophys. **659**, A84 (2022).
[51] W. M. Farr, J. R. Gair, I. Mandel, and C. Cutler, Counting and confusion: Bayesian rate estimation with multiple populations, Phys. Rev. D **91**, 023005 (2015).
[52] B. P. Abbott *et al.*, Supplement: The rate of binary black hole mergers inferred from Advanced LIGO observations surrounding GW150914, Astrophys. J. Suppl. Ser. **227**, 14 (2016).
[53] J. Creighton, Certain identities in FGMC, Technical Report No. T1700029-v2, LIGO DCC, 2017.
[54] M. Dominik, E. Berti, R. O'Shaughnessy, I. Mandel, K. Belczynski, C. Fryer, D. E. Holz, T. Bulik, and F. Pannarale, Double compact objects III: Gravitational wave detection rates, Astrophys. J. **806**, 263 (2015).
[55] I. Harry, J. Calderón Bustillo, and A. Nitz, Searching for the full symphony of black hole binary mergers, Phys. Rev. D **97**, 023004 (2018).
[56] 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).
[57] Y. Pan, A. Buonanno, Y.-b. Chen, and M. Vallisneri, A physical template family for gravitational waves from precessing binaries of spinning compact objects: Application to single spin binaries, Phys. Rev. D **69**, 104017 (2004); **74**, 029905(E) (2006).
[58] 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).