

Flexible and Fast Estimation of Binary Merger Population Distributions with Adaptive KDE

Jam Sadiq ,^{1,*} Thomas Dent ,¹ and Daniel Wysocki ,²

¹*Instituto Galego de Física de Altas Enerxías, Universidade de Santiago de Compostela, Santiago de Compostela, Galicia, Spain*
²*University of Wisconsin–Milwaukee, Milwaukee, Wisconsin, USA*

The LIGO Scientific, Virgo and KAGRA Collaborations recently released the third gravitational wave transient catalog or GWTC-3, significantly expanding the number of gravitational wave (GW) signals. To address the – still uncertain – formation channels of the source compact binaries, their population properties must be characterized. The computational cost of the Bayesian hierarchical methods employed thus far scales with the size of the event catalogs, and such methods have until recently assumed fixed functional forms for the source distribution. Here we propose a fast and flexible method to reconstruct the population of LIGO–Virgo merging black hole (BH) binaries without such assumptions. For sufficiently high event statistics and sufficiently low individual event measurement error (relative to the scale of population features) a kernel density estimator (KDE) reconstruction of the event distribution will be accurate. We improve the accuracy and flexibility of KDE for finite event statistics using an adaptive bandwidth KDE (awKDE). We apply awKDE to publicly released parameter estimates for 44 significant (69) BH binary mergers in GWTC-2 (GWTC-3), in combination with a fast polynomial fit of search sensitivity, to obtain a non-parametric estimate of the mass distribution, and compare to Bayesian hierarchical methods. We also demonstrate a robust peak detection algorithm based on awKDE and use it to calculate the significance of the apparent peak in the BH mass distribution around $35 M_{\odot}$. We find such a peak is very unlikely to have occurred if the true distribution is a featureless power-law (significance of 3.6σ for confident GWTC-2 BBH events, 3.0σ for confident GWTC-3 BBH events).

PACS numbers: 04.30.Tv, 04.30.-w, 04.80.Nn, 97.60.Lf, 98.80.-k, 07.05.Kf, 02.50.-r

Keywords: adaptive width kde — compact binaries

I. INTRODUCTION

Advanced LIGO’s and Advanced Virgo’s [1, 2] first three observing runs have produced many tens of confident detections of binary compact object mergers via gravitational wave (GW) emission, catalogued in GWTC-1 [3], GWTC-2 [4], GWTC-2.1 [5], and GWTC-3 [6].¹ A detailed investigation of the population properties of binary black hole (BBH) mergers, the most commonly detected source type, has been published in [7] focusing on several population characteristics including their component masses and spins using GWTC-2 [4] events. This investigation was updated recently in [8] for GW observations up to the end of the O3 run, as catalogued in GWTC-3 [6], where updated events reported in GWTC-2.1 [5], are included.

One objective has been to reconstruct the primary mass distribution of merging BBH, in order to address questions in stellar evolution, BH formation and binary formation channels. The primary BH masses have lower measurement uncertainty compared to other parameters such as binary spins, thus their distribution is expected to yield significant information for eventual comparison with astrophysical model predictions. The binary chirp mass $\mathcal{M} \equiv (m_1 m_2)^{3/5} (m_1 + m_2)^{-1/5}$, where $m_{1,2}$ are the

source binary masses, is also measurable with (relatively) high precision, thus its distribution is comparable with model predictions stated in terms of binary properties [e.g. 9, 10].

In [7] specific functional forms and Bayesian hierarchical analysis are used to infer the mass distribution of binary mergers from these observed gravitational wave events: see Figure 1 for one output of such inference. These methods are well established, but as our observed sample grows with future observing runs, both computational cost and modeling complexity will increase accordingly. An alternative approach to reconstructing the population is non-parametric methods that allow deviations from a predetermined functional form to better fit observations, or that can approximate arbitrary functional forms. Non-parametric methods implemented via Bayesian analysis have recently been described in [11–17]; the motivation of this work is similar, but we prioritize computational simplicity and speed.

In this paper we propose a fast and flexible method for reconstructing binary merger population distributions via *adaptive bandwidth kernel density estimation* (KDE), using a publicly available code [18]. This method addresses the same goals as the hierarchical analyses mentioned above, but without assuming a particular functional form of the distribution beyond the choice of a kernel function (here, the choice of a Gaussian kernel enforces a smooth estimate). A standard Gaussian KDE with a fixed (global) bandwidth is unlikely to give an accurate representation of the BBH mass distribution, due

* contact: jam.sadiq@usc.es

¹ Each successive catalog update includes previous detections.

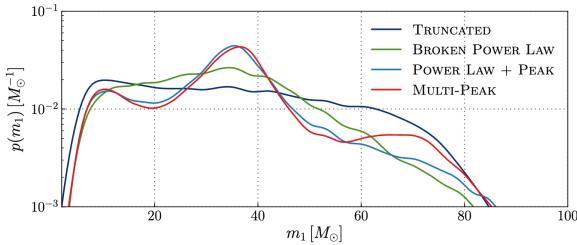


FIG. 1. Observed primary black hole mass distribution predicted by each mass model in [7]; reproduced from that publication.

to its complexity; [19] already identified the risk of significantly biased population estimates due to inappropriate or inflexible bandwidth choice.

Here, we deploy an adaptive KDE (awKDE) estimation method in combination with cross-validation to address the bandwidth problem: specifically, to ensure a locally appropriate bandwidth across a broad parameter space with a wide dynamic range of densities. We estimate the mass distribution of binary mergers from observed events, enabling us to validate and check the model assumptions of the standard Bayesian analyses in a flexible and computationally efficient way; this nonparametric approach can also potentially discover features not described by parametric models. The awKDE estimate of the detected event distribution also yields a differential merger rate estimate, using a fit of the search sensitivity as in [20]. We also compute an uncertainty estimate (confidence interval) for the population KDE using the *bootstrap* technique [21].

There is an apparent peak around $35 M_{\odot}$ in the primary mass distribution of observed GW events, as seen in the POWER LAW + PEAK [22] and MULTI PEAK model results in [7] (see Fig. 1). In this work we also propose a method to determine the significance of one or more prominent peaks in the detected distribution, quantifying the probability that comparable statistical fluctuations might occur in the presence of selection effects. We thus introduce a new peak detection algorithm [23] and apply it to GWTC-2 events. Investigating the presence and significance of such peaks will further help towards understanding BH binary formation; for instance, it has been suggested that pair-instability dynamics in supernovae could induce similar features in the BH mass function [recent discussions include 24–26]; however, as yet no definite signature has been identified.

The paper is organized as follows: In Section II we describe the motivation and reconstruction of the population probability density function using the awKDE algorithm and its application to GW detections, including astrophysical merger rates. We further describe our procedure for computing the uncertainty estimates of our distribution function using bootstrap and awKDE Gaussian kernel contributions. In Section III we apply awKDE to reconstruct the BBH mass distribution from the ob-

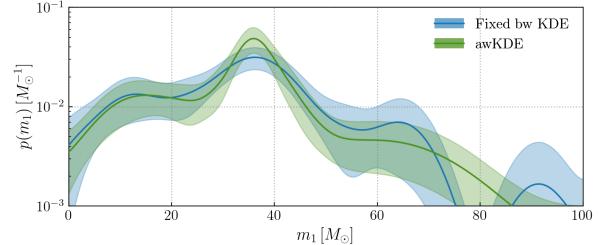


FIG. 2. Fixed global bandwidth versus awKDE reconstruction of primary BH masses using GWTC-2 detected events.

served LVK events. In Section IV we describe a peak detection algorithm to determine the significance of features beyond a power law distribution, and obtain results for GWTC-2 and GWTC-3 detected events. In Section V we summarize our conclusions.

II. RECONSTRUCTION METHOD

We propose the use of a kernel density estimate with adaptive bandwidth selection for reconstructing the probability distribution of source parameters for compact binary mergers observed via GW. This method is nonparametric, straightforward to apply, and enables the identification of general features in the distributions that may be an important input in the astrophysical interpretation of the merging binary population.

While a KDE with fixed (global) bandwidth, chosen for example via Silverman's rule [27], is known to be a suitable method to estimate a distribution close to a single Gaussian, the mass (and possibly also spin) distributions of BBH mergers appear to have a more complex structure, which we do not expect to be well reconstructed by a simple KDE. In particular, the primary and secondary mass distributions may be composed of several components with widely differing mass scales and densities (which have in some cases, e.g. [22], been modelled by power laws and Gaussian peaks). Thus, we consider an extension where the bandwidth varies locally according to an initial estimate of the density of sample points [28, 29]. The differences between a fixed global bandwidth KDE and our proposed awKDE are illustrated in Figure 2.

The fixed bandwidth KDE overestimates the width (underestimates the height) of the observed peak around $35 M_{\odot}$, but also yields a probably unphysical gap in the estimated distribution around $80 M_{\odot}$; these undesirable aspects arise because a fixed bandwidth cannot both reconstruct small-scale features in regions with a high density of points, and supply enough smoothing to avoid artefacts in low density regions.

An adaptive bandwidth KDE [30] is implemented in the open source code AWKDE [18]. We first describe the construction of a KDE from observations X_i , $i = 1 \dots n$: for instance, X_i may be a measured property of a bi-

nary merger. Later, we will describe how measurement uncertainties in these individual event properties are incorporated.

The algorithm computes a density estimate \hat{f} via

$$\hat{f}(x) = n^{-1} \sum_{i=1}^n \frac{1}{h\lambda_i} K\left(\frac{x - X_i}{h\lambda_i}\right), \quad (1)$$

where $K(\cdot)$ is the standard Gaussian kernel,

$$K(z) = \frac{1}{\sqrt{2\pi}} \exp\left(\frac{-z^2}{2}\right), \quad (2)$$

n is the total number of samples and, in general, the product $h\lambda_i$ takes the role of a local bandwidth with h being the global bandwidth.

The first step is the computation of a pilot estimate \hat{f}_0 setting $\lambda_i = 1$ for all i , which is a standard fixed bandwidth KDE. Based on the pilot density \hat{f}_0 , the local bandwidth accounting for variations in the density of samples is obtained via

$$\lambda_i = \left(\frac{\hat{f}_0(X_i)}{g}\right)^{-\alpha}, \quad (3)$$

where α is the local bandwidth sensitivity parameter ($0 < \alpha \leq 1$) and g is a normalization factor

$$\log g = n^{-1} \sum_{i=1}^n \log \hat{f}_0(X_i). \quad (4)$$

Finally the adaptive KDE $\hat{f}(x)$ is obtained by evaluating (1) with the variable (local) bandwidth $h\lambda_i$. The above method assumes one-dimensional data X_i ; the method may also be applied to two- or more-dimensional data by linearly transforming the data to have zero mean and unit covariance.

The method requires a choice of the initial global bandwidth h and sensitivity parameter α : we use the *leave-one-out cross-validation* method [31] to determine these values. As a figure of merit for the cross-validation we use the (log) likelihood,

$$\log \mathcal{L}_{\text{LOO}} = \sum_{i=1}^n \log \hat{f}_{\text{LOO},i}(X_i), \quad (5)$$

where $\hat{f}_{\text{LOO},i}$ is the KDE constructed from all samples *except* X_i . This choice, as it is linear in the logarithm of the estimate at observed values, will penalize *relative* errors. Since we wish to obtain an accurate estimate of densities over a large dynamic range, the log likelihood is more suitable than FOM based on absolute error or squared absolute error. We employ a grid search over a range of h and α values; however, we often find the likelihood is maximized with α at or close to 1, thus in some cases (discussed later) we will impose $\alpha = 1$ rather than conduct the full 2d grid search.

A. Application to GW observations

The component masses of observed GW binary mergers have significant measurement uncertainties, which we wish to incorporate in reconstructions of the mass distribution. The parameters of each binary are obtained by Bayesian inference using models of the emitted GW waveform [e.g. 32], referred to as parameter estimation (PE). We consider detected mergers labelled by $i = 1 \dots n$: uncertainties in a given parameter X are quantified via random values X_i^k drawn from the posterior of the i th merger. Typically, some thousands of parameter samples are available per merger event [33].²

Although we use a random selection of PE samples for the KDE including measurement uncertainty, to first obtain the optimum global bandwidth and sensitivity parameter α we evaluate the likelihood \mathcal{L}_{LOO} using only the *median* parameter value for each merger as the data X_i . This choice reduces the computational cost and avoids over-fitting of random fluctuations: note that these medians are independent values, whereas the PE samples for a given event i , X_i^k , are not independent of one other, as they are all correlated with the (unknown) true source mass: thus, the PE samples for a given event are not independent draws from the population distribution that the KDE aims to reconstruct.

Having obtained the optimal h and α choices, we construct the population KDE using 100 randomly chosen samples for each event. We verify that increasing the number of PE samples did not change our results significantly: considering KDE outputs using 10, 30, 100, and 300 PE samples and taking the 300-sample fit as a baseline, we find the maximum fractional difference, $\max_{m_1} f_{300}(m_1)^{-1} [f_n(m_1) - f_{300}(m_1)]$, for the 10-, 30-, and 100-sample outputs to be $\sim 20\%$, $\sim 15\%$, and $\sim 6\%$, respectively. We note though that for large measurement errors relative to the scale of structures in the underlying distribution, the KDE result is likely to be over-dispersed, i.e., any sudden variations in the actual density will be smoothed out.³

Our major source of uncertainty for the population KDE lies in the finite number of binary merger events, and the resulting count fluctuations in the density estimate at a given parameter value. We estimate this uncertainty by bootstrap resampling [21] over the merger events i : i.e., we create a large number of bootstrapped data sets, each containing some number of copies of the PE samples for each event, with the number of copies following a binomial distribution. For the cases we consider, this distribution is well approximated by the Poisson distribution with unit mean, hence we expect this procedure

² Note that we use samples obtained with a distance prior proportional to the comoving volume element times (redshifted) source time.

³ In principle this bias can be tackled by the computationally demanding hierarchical methods, if they model functional forms that are sufficiently close to the true distribution.

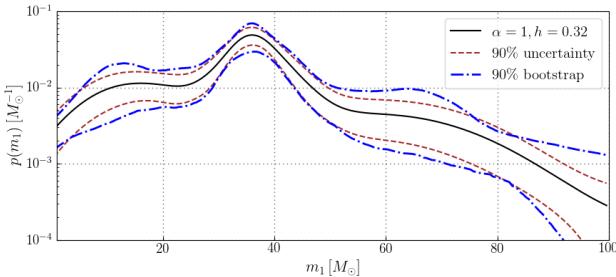


FIG. 3. awKDE (black solid line) for GWTC-2 events with error estimates using a standard bootstrap method (blue dot-dashed lines) and from the Gaussian Kernel coefficients from awKDE code (brown dashed lines) using Eq. (8). The two error estimates are in good agreement for most parameter values.

to give the correct scaling of uncertainties due to event count fluctuations. We extract the median and 5th and 95th percentiles from 1000 bootstrap iterations to obtain the confidence interval at any given parameter value; by plotting all bootstrap KDEs we can also visually identify regions of high and low uncertainty.

B. Fast approximate uncertainty estimate

While constructing an awKDE with median parameter values from PE, or in general using one parameter value per event, we can obtain a rapid uncertainty estimate by considering the awKDE Gaussian kernel contributions due to observed events at any given parameter value x (accessible with our modified awKDE algorithm [18]). These contributions or KDE coefficients are

$$c_k(x) = n^{-1} \frac{1}{h\lambda_k} K\left(\frac{x - X_k}{h\lambda_k}\right), \quad (6)$$

where k labels the observed events. We may estimate the standard deviation of the coefficients, σ_c , as

$$\sigma_c^2(x) \simeq \langle c_k^2 \rangle - \langle c_k \rangle^2, \quad (7)$$

where $\langle \cdot \rangle$ represents the mean over observed events. As the KDE is the sum of n such coefficients its variance is a factor n larger, thus we take $\sigma_{\text{KDE}}(x) = \sqrt{n}\sigma_c(x)$. The resulting estimated 90% confidence interval is $1.64\sigma_{\text{KDE}}(x)$ above and below the central value,⁴ i.e.

$$\epsilon_{\text{KDE}}(x) = 1.64\sqrt{n}\sigma_c(x). \quad (8)$$

We compare this estimate with the bootstrap uncertainty region and find them in a good agreement for most parameter values, as shown in Figure 3. Clearly this direct

variance estimate is computationally much more efficient than the bootstrap.

While this coefficient-based estimate agrees with the bootstrap for small bandwidths (relative to the range of sample values X_i), for large values of bandwidth it becomes significantly smaller than the bootstrap uncertainty. We also see from Fig. 3 that the coefficient-based estimate tends to underestimate the bootstrap uncertainty in regions of low density. A complete explanation of this behaviour should include more complicated effects due to the two-step nature of the adaptive KDE: for instance, the estimate of (8) does not account for any systematic variation in coefficients due to changes in the bandwidths λ_k between different data realizations. Also, the coefficient-based estimate considers only the variance of the density, which does not account for uncertainty distributions that deviate significantly from Gaussian, such as Poisson statistics in the small number regime.

We may get some insight from heuristic arguments: if we approximate the Gaussian kernel coefficients at a given x as constants c_L ('large' coefficients) for n_L data points X_k 'close' to x relative to their bandwidths $h\lambda_k$, and vanishing for the remaining $n - n_L$ points 'far away' from x , then $\sigma_{\text{KDE}}(x)$ becomes approximately $(c_L n_L/n) \sqrt{n_L^{-1} - n^{-1}}$. This corresponds to the variance of a binomial distribution with a success rate n_L/n . When the bandwidths are large, n_L will be close to n (many events will have large and near-equal coefficients), leading to a cancellation.

However, given that the data are produced by an underlying Poisson process, as the total number of astrophysical events observed is not fixed, we also considered a modified estimate which scales with the variance of the count of 'close' events, reducing to $(c_L n_L/n) \sqrt{n_L^{-1}}$. This corresponds to removing the second term in Eq. (7), i.e.

$$\hat{\sigma}_c^2(x) = \langle c_k^2 \rangle. \quad (9)$$

We find that the 90% error $\hat{\epsilon}_{\text{KDE}}(x)$ corresponding to this modified estimate agrees much more closely with bootstrap uncertainties for large bandwidths. This property motivates our use of $\hat{\epsilon}_{\text{KDE}}$ in the peak detection method to be introduced in Section IV.

III. AWKDE POPULATION RECONSTRUCTION FROM LIGO-VIRGO DETECTIONS

In this section we apply the methods described above to reconstruct the distribution of parameters for BBH observed so far in the advanced detector runs. We divide this section into results obtained using the GWTC-2 event set, comprising events from the O1, O2 and O3a runs [3, 4], which we compare to the collaboration results in Abbott *et al.* [7]; and results using the GWTC-3 set [6] (including an updated release for data and events up to O3a, called GWTC-2.1 [5], where available).

⁴ The factor 1.64 corresponds to the 95th percentile of the standard normal distribution.

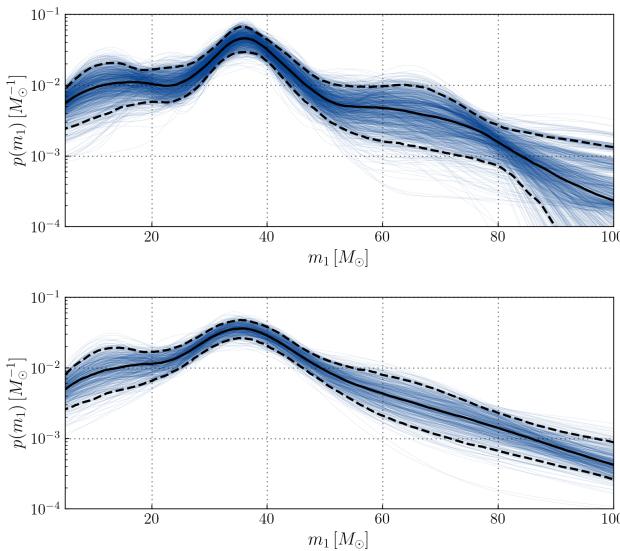


FIG. 4. KDE with uncertainty estimates of source frame primary mass, m_1 , for detected BBH events in the O1, O2 and O3a observing runs. Top: KDE using the PE sample median for each event. Bottom: KDE constructed using 100 random samples from each observed event. The median (black solid), 90% confidence interval (black dashed lines) and blue curves are constructed using a bootstrap. Our KDE results match closely the POWER LAW + PEAK [22] and MULTI PEAK model results in [7], see Fig. 1.

A. BBH population from GWTC-2

We use detected GW events from GWTC-2 [3, 4] to reconstruct the KDE of parameters for BBH observations. As in [7], we consider significant events (false alarm rate below 1/yr) for which both component masses are above $3 M_\odot$, finding 44 such mergers. We first consider the primary mass m_1 defined in the binary source frame. The resulting KDE bootstrap samples, median estimate and 90% confidence intervals are shown in Figure 4, where the top panel uses only the median m_1 value for each binary merger and the lower panel accounts for measurement uncertainties via the PE samples. As described in Sec. II A, the KDE bandwidth and sensitivity parameter α are optimized via leave-one-out cross validation using median m_1 sample values, and the same optimized parameters are used to evaluate the KDE including PE sampling of mass uncertainties.

The resulting primary mass distributions match well those inferred using parametric models in [7], specifically with the POWER LAW + PEAK [22] and MULTI PEAK models most preferred by the GWTC-2 data. We note that there is a clear global maximum in the observed distribution at $\sim 35 M_\odot$ and that the uncertainty in the mass distribution is smallest here. The awKDE is also able to reconstruct a wide dynamic range of densities without excessive statistical uncertainties.

As an alternative to the primary mass, we also consider

the chirp mass \mathcal{M} for observed BBH mergers, which is also measured with relatively small uncertainty. The corresponding KDEs are shown in Figure 5. Here in addition

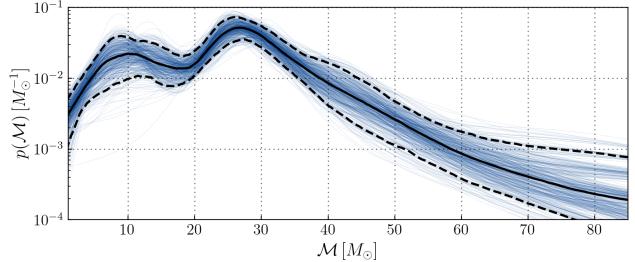


FIG. 5. KDE with uncertainty estimates of source chirp mass, \mathcal{M} , for detected events in the GWTC-2 or O1, O2 and O3a observing runs. KDE constructed using 100 random samples from each observed event. The median (black solid), 90% confidence interval (black dashed lines) and blue curves are constructed using a bootstrap. In addition to the principal peak in the distribution around $25-30 M_\odot$, there is a hint of further structure around $10 M_\odot$.

to the expected peak just below $30 M_\odot$ we note hints of a secondary peak around $\mathcal{M} \simeq 10$, though small compared to the estimated uncertainties: however, with this data we are not able to distinguish the multiple peaks claimed in [11] from random fluctuations.

a. Merger rate estimation from awKDE We also consider estimation of BBH merger rates using our awKDE results. The additional ingredient in this analysis is the sensitive volume-time VT over which a source is detectable: we quantify this based on an approximation for the sensitivity of GW detectors, with corrections to account for the actual behaviour of searches in real data [20]. The idea is to calibrate a semi-analytic function VT_{analytic} against the re-weighted result of injections (i.e. simulated signals added to real data and analyzed by search pipelines) VT_{inj} , assuming a parameterized relationship between them that can be expressed via basis functions.

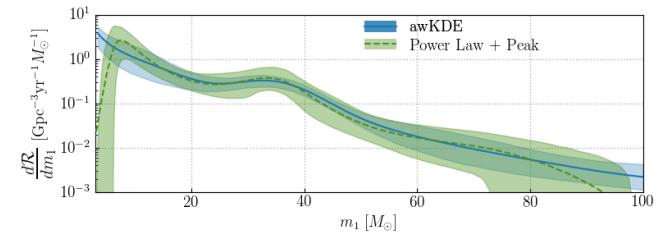


FIG. 6. Rate estimates using adaptive KDE and sensitive volume for BBH events in GWTC-2. The green band curve is the rate estimate using the POWER LAW + PEAK model from [7]; the blue band shows a rate estimate using KDE results from median PE values and a fit to search sensitivity.

For the GWTC-2 observing runs, the ‘uncalibrated’ semi-analytic sensitivity is estimated using the criterion

that the signal-to-noise ratio (SNR) of a signal in the second most sensitive detector, using specific reference power spectral densities, should be greater than 8. This semi-analytic estimate VT_{analytic} is obtained for a given observing time T over a grid of intrinsic source parameters such as component masses. For the injection VT_{inj} , one performs a set of injections and counts the number detected by search pipelines to obtain an average VT_{inj} for the injected population at given intrinsic parameters. The calibrated estimate of VT is then obtained by fitting a correction function to the deviations of VT_{inj} from VT_{analytic} , as described in [34] and applied in [7].

We compute the merger rate density using this corrected VT and our KDE results for the primary mass $\hat{f}(m_1)$ from observations up to O3a. We obtain $VT(m_1)$ assuming a power law distribution for the secondary mass m_2 , taking the median power $\beta_q = 1.26$ from the TRUNCATED model samples from [7], and find the merger rate density via

$$\frac{dR}{dm_1} = n \frac{\hat{f}(m_1)}{VT(m_1)}. \quad (10)$$

Our estimate is in good agreement with the results in [7] for the POWER LAW + PEAK model, as shown in Fig. 6.

B. BBH population from GWTC-3

We also used recently available public data from GWTC-3 [6] to construct the distribution of primary and chirp masses using PE samples for 69 confident BBH events with a false alarm rate below 0.25/yr. Note that this data set includes a small number of events identified in a reanalysis of O3a data with improved search pipelines, detailed in GWTC-2.1 [5]. We find new features in the distributions, consistent with the results in [8]. As with the GWTC-2 data, we do not recover significant structure beyond the two visible peaks. The peak at $35 - 40 M_\odot$ in primary mass ($25 - 30 M_\odot$ in chirp mass) shows signs of asymmetry (skewness) relative to a locally Gaussian form, which may indicate a need for more complex models.

a. Merger rate estimation from awKDE We also estimate BBH differential merger rate over m_1 using our awKDE results from GWTC-3 data and an estimate of sensitive volume-time VT from GWTC-3 searches, and compare the estimate with the results of [8], as shown in Figure 8. Here we see similar features to those reported by the FLEXIBLE MIXTURES, POWER LAW + PEAK, and POWER LAW + SPLINE models in [8]. The structure of the rate distribution obtained from KDE does differ in detail from these hierarchical analyses, probably due to our different treatment of source mass measurement uncertainties, and their use of a specific model assumption for the binary mass ratio q (i.e. a power law distribution). We also estimated dR/dm_1 by making a two-dimensional KDE reconstruction of detected events over m_1, m_2 , converting to a rate distribution by dividing by $VT(m_1, m_2)$

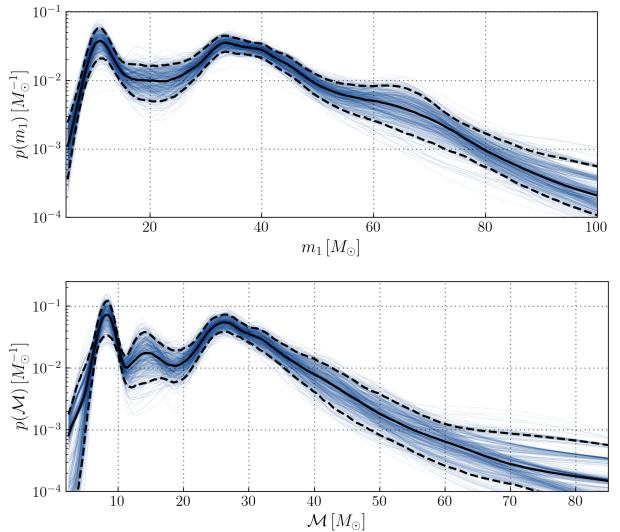


FIG. 7. Application of awKDE to GWTC-3 data [6]: we reconstruct the source primary mass m_1 and chirp mass \mathcal{M} distributions, including measurement uncertainties via PE samples, seeing additional features at lower masses consistent with the results in [8]. (The apparently multi-modal estimate at high \mathcal{M} is a by-product of the bootstrap method in regions with very low event count.)

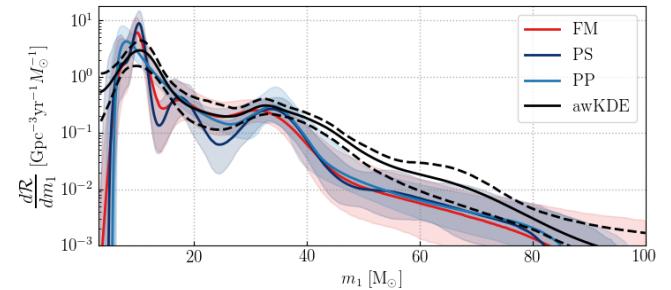


FIG. 8. Rate estimates using adaptive KDE and sensitive volume for BBH events in GWTC-3. Black curves show the rate estimate using the KDE reconstruction in the top panel of Fig. 7 and a fit to search sensitivity, compared with the FLEXIBLE MIXTURES (FM), POWER LAW + SPLINE (PS), and POWER LAW + PEAK (PP) mass models, as reported in [8].

(obviating any assumption on the distribution of q), and marginalizing over m_2 : the resulting rate is consistent with our estimate in Fig. 8.

C. 2D KDE for cosmological evolution of BH mass distribution

We apply our adaptive KDE in a two-dimensional space using median PE values for m_1 in the source frame and luminosity distance D_L for BBH events in GWTC-3 [6], following the same procedure of leave-one-out cross-validation to determine the optimal global bandwidth

and sensitivity parameter. As shown in Figure 9, the resulting two-dimensional KDE shows peaks around $10M_{\odot}$ and $35M_{\odot}$, similar to the one-dimensional mass KDEs and consistent with results in [8]; these overdensities appear to be present consistently over different distances. The absence of detected events at low mass and higher D_L is due to the selection function, i.e. the dependence of search sensitivity on mass.

We also show the relative uncertainty in the estimated density using bootstrap method, see bottom panel of Figure 9, showing that over almost all the region populated by detections the uncertainty is significantly less than the density; hence the apparent peaks are unlikely to be entirely due to random counting fluctuations.

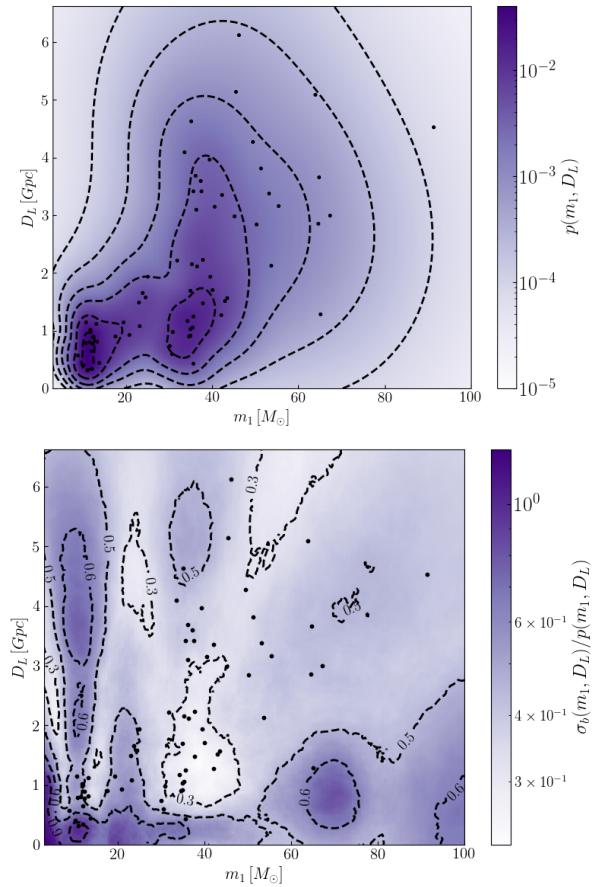


FIG. 9. Adaptive KDE over the source frame primary mass and luminosity distance for BBH events in GWTC-3. The first plot shows the distribution reconstructed from 100 PE samples per event. Features around $10M_{\odot}$ and $35M_{\odot}$ appear to be present consistently over different distances. The second plot shows a relative uncertainty estimate: $\sigma_b(m_1, D_L)/p(m_1, D_L)$, where $\sigma_b = (\sigma^+ + \sigma^-)/2$, the average of upper and lower 1- σ bootstrap error estimates using 100 PE samples per event. The relative error is < 1 for most regions populated by events.

IV. DETECTION OF BH MASS FUNCTION PEAKS WITH AWKDE

In this section we describe our algorithm [23] to detect prominent peaks in the BH mass function with awKDE. Peaks in the mass spectrum can provide hints of the formation channels responsible for the observed BBH (e.g. [25]). As shown in Figure 1, the two parameterised models of [7] that allow for a Gaussian peak component find such a feature around $35M_{\odot}$, as also visible in the awKDE reconstruction. However, the statistical preference for such models over power-law based models with no peak component or other feature is only moderate (log10 Bayes factor < 2 for GWTC-2 events). Here we pursue an alternative strategy by developing a general technique to identify peaks in the detected mass function, and a procedure to estimate their statistical significance.

A. Peak Detection Algorithm

We first propose an algorithm for detection of the most prominent peak in a (differentiable) one-dimensional distribution, which we then apply to our adaptive width KDEs. We consider the following algorithm:

- As a pre-processing step, we cancel the maximum likelihood power-law dependence from the density estimate, in order to obtain a distribution that is as close to uniform as possible in the absence of peaks: see IV C below.
- For a given distribution $\hat{f}(x)$, find all local maxima, denoted $\{x_j^P\}$.⁵
- Using a window of given size δ (the choice of δ is discussed below), evaluate the PDF on either side of the peak, $\hat{f}_j^{\pm} = \hat{f}(x_j^P \pm \delta)$. In the event that $x_k^P \pm \delta$ is outside the range where $\hat{f}(x)$ is defined, to avoid extrapolation we substitute \hat{f}_j^+ (\hat{f}_j^-) with the value of \hat{f} at the highest (lowest) x_j in the available range.
- Compute peak heights using

$$H_j^P = \hat{f}(x_j^P) - (\hat{f}_j^- + \hat{f}_j^+)/2. \quad (11)$$

- Evaluate the estimated uncertainty at each peak, ϵ_j^P , in our case defined in Eq. (8).⁶

⁵ In practice the algorithm requires the number of local maxima, determined numerically, to not be very large, which imposes some constraint on the smoothness of the distribution.

⁶ If an uncertainty estimate is not available, only H_j^P can be used to define the detection statistic in the final step.

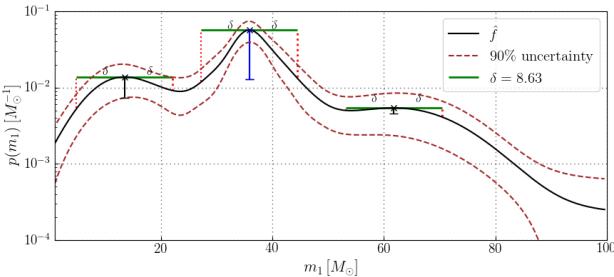


FIG. 10. Peak detection algorithm for KDEs applied to BBH primary masses in GWTC-2 with an arbitrary choice of global bandwidth h . The solid black curve is the awKDE estimate \hat{f} with the dashed brown curves indicating 5th and 95th percentile uncertainties estimated via Eq. (8). The green lines indicate the window size δ , vertical red dotted lines indicate the construction used to measure peak height, while blue (black) vertical lines indicate the most significant (less significant) peak heights. (For simplicity, here we do not illustrate the pre-processing step to cancel the maximum likelihood power law dependence.)

- Determine the most significant peak by maximizing a detection statistic which combines H_j^p and ϵ_j^p (we discuss the choice of detection statistic below).

This maximized detection statistic value will be used to distinguish distributions with significant peak features from those without such features. Figure 10 illustrates the four steps of the algorithm subsequent to the post-processing step, applied to a set of BBH primary masses.

B. Mock data sets and algorithm evaluation

In order to evaluate and optimize the detection performance of our algorithm we created several large-scale mock datasets, each having 10,000 samples with 60 data points in each sample. Using the detection statistic values produced by the algorithm for background and signal samples, we calculate receiver operating characteristic curve (ROCs) to measure the efficiency of peak detection. Here, the false alarm or false positive probabilities are defined under the assumption of a truncated power-law population of mergers (accounting for the detection selection function), which we consider as representing a featureless distribution with no peak(s).

To represent the signal hypothesis we construct a mock dataset referred to as the Peak dataset, with samples containing a mixture of a uniform (random) component and a Gaussian component. The fraction of points drawn from the Gaussian distribution is chosen randomly between 5%–95%, the uniform component is taken to lie between $3 M_\odot$ – $100 M_\odot$, while the mean of the Gaussian component is uniformly distributed between $8 M_\odot$ – $51.5 M_\odot$ and its standard deviation uniformly distributed on $5 M_\odot$ – $10 M_\odot$.

For the background, i.e. no-peak hypothesis, we con-

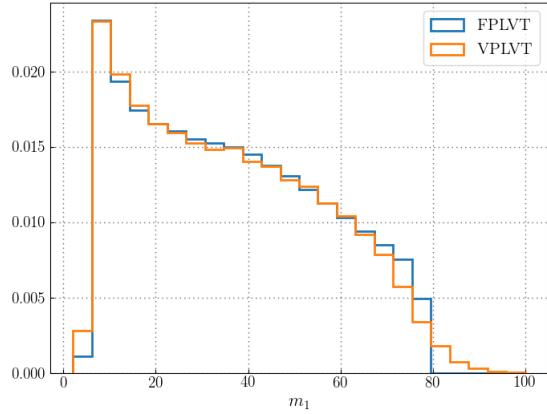


FIG. 11. Distribution of background data with fixed/variable hyperparameter power law VT corrected (FPLVT/ VPLVT) datasets, based on GWTC-2 population analysis with a truncated power law model.

struct three datasets, also with 60 data points per sample, that more or less closely represent possible detected events if the true astrophysical distribution is a featureless power law. The first Uniform dataset contains only uniform random samples, the simplest realisation of a null case. The two other background datasets are drawn from a power law distribution of BBH primary and secondary masses, the TRUNCATED mass model in [7] (“Model B” in [35], see references therein). The model hyperparameters were estimated via Monte Carlo sampling: in our two datasets we either consider fixed hyperparameter values given by the sample medians, or variable hyperparameter values using available publicly released samples. With fixed parameter values, the primary mass power-law spectral index α is set to 2.21, the lower cut-off m_{\min} to 5.97, the upper cut-off m_{\max} to 78.47, and the mass ratio $q \equiv m_2/m_1$ follows a power-law with spectral index β set to 1.26.

After generating a large number of m_1 , m_2 samples from the TRUNCATED model we account for the selection function. The probability of detection over O1, O2 and O3a data is proportional to the search sensitive volume \times time (VT), estimated as a function of (m_1, m_2) via a semi-analytic calculation calibrated to search pipeline injection results [34] as in Section III A 0 a. Given the maximum value of VT over masses, we use rejection sampling to implement the selection function. For the fixed hyperparameters case we obtain 60,000 data points after rejection sampling, yielding 10,000 independent samples with 60 points each. For the varying hyperparameters case we obtain 60 points for each hyperparameter sample; 7460 such samples are available from the analysis of Abbott *et al.* [7]. For GWTC-3, hyperparameter samples for the TRUNCATED model are not publicly available, however we were able to obtain so far

unpublished results from which we use 10,000 samples,⁷ following the same procedure as for GWTC-2.

The resulting fixed/variable hyperparameter power-law VT corrected datasets are denoted FPLVT/ VPLVT. The distribution of primary masses in these power-law based datasets is shown in Fig. 11 for the GWTC-2 case. This distribution is clearly far from uniform, however it does not have a well-defined local maximum feature apart from the cutoff at low mass.

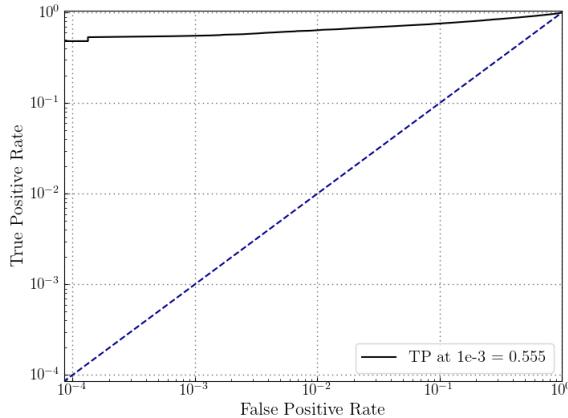


FIG. 12. ROC curve using the VPLVT dataset for GWTC-2 based power law samples as background and the Peak dataset as signal, for an arbitrary choice of detection statistic. In this case the detection algorithm gives 55.5% true positive rate at a false alarm probability of 10^{-3} .

For any specific choice of algorithm, we evaluate detection statistic values for the Peak dataset and for each of the three background datasets and then compute ROCs as shown in Figure 12. Our figure of merit is the true positive rate at a false positive rate or false alarm probability (FAP) of 10^{-3} . We use this criterion to tune the fixed window size δ the choice of global KDE bandwidth and the peak detection statistic.

C. Optimisation of peak detection algorithm

As we see in Figure 11, there is a power law-like dependence in the mock datasets which can cause the peak detection algorithm to identify spurious peaks at small m_1 values. As the algorithm is designed to identify peaks above a background of nearly uniform density, in order to improve its sensitivity we apply a pre-processing step that cancels the maximum likelihood power-law dependence from the density estimate. With the power-law

probability density given by

$$f(m) = \frac{m^\gamma}{\int_{m_{\min}}^{m_{\max}} dm m^\gamma}, \quad (12)$$

the log likelihood for N events with mass values m_i is

$$\ln \Lambda = \sum_{i=1}^N \ln \frac{m_i^\gamma}{\int_{m_{\min}}^{m_{\max}} dm m^\gamma}. \quad (13)$$

We obtain the maximum likelihood value γ_{ML} value setting $d \ln \Lambda / d\gamma = 0$ and numerically solving the resulting equation. Then, our pre-processing step multiplies the KDE by $m^{-\gamma_{\text{ML}}}$ in order to suppress spurious peaks arising from a (non-uniform) power-law underlying distribution.

D. Tuning the peak window δ

In our algorithm we need to choose a window size δ to compute peak heights. We considered both fixed (pre-determined) δ choices, and values determined by the data or by the KDE construction.

The highest figures of merit were obtained by setting $\delta = kh$, where h is the standard deviation of the Gaussian kernel used to construct the initial pilot (fixed bandwidth) KDE, i.e. the global bandwidth, and k is a constant factor which we took to range from 1 to 8. Depending on other choices, we will eventually identify the value of k with the highest FOM. Given our method of selecting h , the choice $\delta = kh$ ensures that the window size is automatically scaled to the variance of the data, whereas any fixed δ value would have to be retuned when applying the method to other physical situations.

E. Tuning the global KDE bandwidth

In order to choose the global bandwidth for the pilot KDE we do not consider leave-one-out cross validation, as for the 10,000 samples of our background data this would be too computationally expensive. Also, the optimal bandwidth for accurate reconstruction of the mass distribution is not necessarily most suited to detection of peaks. Instead, we construct an awKDE for each sample, fixing $\alpha = 1$, for several different choices of bandwidth between 0.1 and 0.5 times the standard deviation of the data points, uniformly spaced in $\log(h)$. We restrict the maximum bandwidth to 0.5 in units of the data standard deviation since if the bandwidth is larger the resulting KDE is necessarily very smooth, without smaller scale features. This affects the goal of our algorithm to identify datasets which have a peak feature smaller than the total overall extent of the distribution, versus those without such a feature. Out of these different bandwidths for a given sample, we select the one for which the maximum peak detection statistic is highest.

⁷ A. Farah, private communication.

F. Range of KDE evaluation and boundary conditions

In order to check any specific trends in our algorithm we plotted several diagnostic tests including peak height, peak detection statistic and peak location for our mock datasets. We found that some spurious peaks were generated in peak datasets due to the falloff of the KDE towards the edges of the range, a known artifact resulting from the absence of points outside the range [36, 37]. To suppress these spurious peaks we use a bounded KDE, effectively reflecting the data points about the edges of the range, and we check that this choice increases the detection figure of merit. The choice of range of values over which to evaluate the KDE also affects the detection algorithm: we find the best performance for a KDE range which extends the range of the data points by 10% below (above) the minimum (maximum) points in a given sample.

G. Tuning choice of peak detection statistic

We lastly need to specify the detection statistic which determines the most significant peak in each awKDE. We considered various possibilities: the simplest choice would be to use the peak height H_j^P , however we also investigated combinations of the peak height and the error estimate $\epsilon_{\text{KDE}}(x_j^P)$ at the peak as in Eq. (8). Considering the ratio $H_j^P/\epsilon_{\text{KDE}}(x_j^P)$ as detection statistic, it has higher FOM than simply the peak height. However, the error estimate $\epsilon_{\text{KDE}}(x_j^P)$ reduces sharply with increasing bandwidth, leading to some spurious peaks in our background datasets with small height but smaller estimated errors.

Finally, to avoid such artefacts at high bandwidth we consider the modified uncertainty estimate of Eq. (9), which better agrees with the uncertainty obtained from bootstrap resampling. Using the statistic $H_j^P/\hat{\epsilon}_{\text{KDE}}(x_j^P)$ the FOM increased further, thus this expression is our final choice.

H. Optimized detection choice

After the tuning steps described above, our choice of peak detection statistic is the ratio $H_j^P/\hat{\epsilon}_{\text{KDE}}(x_j^P)$: then for the FPLVT background dataset $\delta = 3h$ gives the highest FOM, whereas for the VPLVT dataset $\delta = 5h$ gives the highest FOM (nearby choices give very similar performance). Hence we use $\delta = 4h$ to compute the detection statistic for the real dataset consisting of observed GW signals from BBH merger.

I. Peak detection for GWTC-2 BBH events

We apply our optimized detection algorithm to the 44 high-significance BBH detections in GWTC-2 as used in [7], finding a peak detection statistic of 2.85. To assess the significance of this value we construct comparable background samples, each having 44 data points, analogously to the FPLVT (VPLVT) mock data described in the previous section; as before we obtain 10000 (7640) samples and apply the same algorithm to compute their detection statistics.

We then compute the FAP for the detected events up to O3a using these background datasets, obtaining a value 0.0001 (3.7σ significance) for the FPLVT background and a value of 0.00013 (3.6σ) for the VPLVT background.⁸

So far, we have not accounted for mass measurement error in constructing background samples. We expect that such errors would not cause more significant peaks to appear in the KDE of BBH primary masses, however we investigate this by simulating errors of similar magnitude to those in the primary masses of GWTC-2 events [4]. We approximate the distribution of errors in m_1 as log-normal, and find a representative standard deviation of 0.24 by considering the published 90% credible intervals. Applying errors with this distribution to the VPLVT background dataset, we recompute the detection statistic values and obtain a FAP for the detected events up to O3a of 0.0017 (2.9σ). Hence such errors do not strongly impact our conclusions.

Thus, the apparent feature in GWTC-2 BBH primary masses is very unlikely to have originated from random fluctuations, if the true underlying distribution was a power law. This result strengthens the motivation to search for possible astrophysical explanations of such features.

J. Peak detection for GWTC-3 BBH events

We further applied the detection algorithm to 69 high-significance (false alarm rate below 0.25/yr) BBH events in GWTC-3 [6], finding a peak detection statistic of 3.36. Using FPLVT (VPLVT) mock data with 10000 samples, each having 69 data points, we compute their detection statistics. We then compute the FAP for the detected events in GWTC-3 using these background datasets, obtaining values 0.0001 (3.7σ) for the FPLVT background and 0.0012 (3.0σ) for the VPLVT background. Following the same procedure as for GWTC-2 to approximate the distribution of errors in m_1 as log-normal, we find a representative standard deviation of 0.23 by considering the published 90% credible intervals. Applying errors

⁸ These results are obtained using the pre-processing procedure of Sec.IV C; in our previous analysis without this step we obtained FAP values 0.0005 (3.3σ) for the FPLVT background and 0.0016 (2.9σ) for VPLVT background.

with this distribution to the VPLVT background dataset, we recompute the detection statistic values and obtain a FAP for GWTC-3 events of 0.0103 (2.3 σ).

This apparently reduced significance for GWTC-3 events against VPLVT backgrounds, despite the higher peak detection statistic, reflects a difference in the TRUNCATED hyperparameter samples used for the background: specifically a more negative power-law index for the primary mass in GWTC-3 as compared to GWTC-2, leading to spurious peaks with somewhat higher detection statistics despite the steps taken in our algorithm. Thus, it is likely that the algorithm could be further refined to better reject steep power law background distributions. Our astrophysical conclusion that the primary BH mass distribution contains a peak-like feature strongly inconsistent with a power law is not affected.

V. CONCLUSIONS

GW observations from compact binary mergers detected in LVK data provide key astrophysical constraints, by allowing the comparison of theoretical models to the observed population of BBH. Many techniques to measure this population rely on a predetermined functional model form (e.g., [38], [22]), which cannot react to surprises without human intervention. In this work we provide a new, computationally cheap method for measuring the population with weak assumptions, allowing for the automated discovery of new population features; it will thus help to answer questions regarding mass gaps, sub-populations, and other signatures of BH formation channels as the number of detections increases.

Our fast and transparent method uses adaptive bandwidth KDE (awKDE) [18] to estimate or reconstruct the population distribution of compact binary mergers detectable by the global gravitational-wave interferometer network. The resulting estimates can serve as sanity checks on computationally more expensive Bayesian hierarchical inferences on binary merger populations, and may also be able to bring to light features that are not present in commonly used parameterized models. The method can straightforwardly incorporate uncertainties in individual source parameters and yield estimates of astrophysical merger rates. We specifically demonstrated its use for binary black hole (BBH) merger distributions over primary mass, chirp mass and cosmological distance.

In addition, we introduce a robust peak detection algorithm [23] for smooth one-dimensional distributions, such

as those produced by awKDE, and apply the algorithm to BBH primary masses from the second LIGO-Virgo catalog of transient GW sources, obtaining significant evidence for the presence of a peak feature in the mass distribution as opposed to a simple power law. Possible future developments of this algorithm include extensions to detect gaps or dips in the distribution, to detect multiple features, or to better account for dependence of the detection probability on binary masses (or other parameters).

ACKNOWLEDGEMENTS

The authors thank Amanda Farah for sharing unpublished population analysis results for the TRUNCATED model using GWTC-3 data. This work has received financial support from Xunta de Galicia (Centro singular de investigación de Galicia accreditation 2019-2022), by European Union ERDF, and by the “María de Maeztu” Units of Excellence program MDM-2016-0692 and the Spanish Research State Agency. DW thanks the NSF (PHY-1912649) for support. This material is based upon work supported by NSF’s LIGO Laboratory which is a major facility fully funded by the National Science Foundation. The authors are grateful for computational resources provided by the LIGO Laboratory and supported by National Science Foundation Grants PHY-0757058 and PHY-0823459. This research has made use of data, software and/or web tools obtained from the Gravitational Wave Open Science Center (<https://www.gw-openscience.org/>), a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. 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.

-
- [1] J. Aasi *et al.* (LIGO Scientific), *Class. Quant. Grav.* **32**, 074001 (2015), [arXiv:1411.4547 \[gr-qc\]](https://arxiv.org/abs/1411.4547).
- [2] F. Acernese *et al.* (VIRGO), *Class. Quant. Grav.* **32**, 024001 (2015), [arXiv:1408.3978 \[gr-qc\]](https://arxiv.org/abs/1408.3978).
- [3] B. P. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. X* **9**, 031040 (2019), [arXiv:1811.12907 \[astro-ph.HE\]](https://arxiv.org/abs/1811.12907).
- [4] R. Abbott *et al.* (LIGO Scientific, Virgo), *Phys. Rev. X* **11**, 021053 (2021), [arXiv:2010.14527 \[gr-qc\]](https://arxiv.org/abs/2010.14527).
- [5] R. Abbott *et al.* (LIGO Scientific, VIRGO), (2021), [arXiv:2108.01045 \[gr-qc\]](https://arxiv.org/abs/2108.01045).
- [6] R. Abbott *et al.* (LIGO Scientific, VIRGO, KAGRA), (2021), [arXiv:2111.03606 \[gr-qc\]](https://arxiv.org/abs/2111.03606).

- [7] R. Abbott *et al.* (LIGO Scientific, Virgo), (2020), [arXiv:2010.14533 \[astro-ph.HE\]](https://arxiv.org/abs/2010.14533).
- [8] R. Abbott *et al.* (LIGO Scientific, VIRGO, KAGRA), (2021), [arXiv:2111.03634 \[astro-ph.HE\]](https://arxiv.org/abs/2111.03634).
- [9] F. S. Broekgaarden *et al.*, (2021), [arXiv:2112.05763 \[astro-ph.HE\]](https://arxiv.org/abs/2112.05763).
- [10] M. Dominik, E. Berti, R. O’Shaughnessy, I. Mandel, K. Belczynski, C. Fryer, D. E. Holz, T. Bulik, and F. Pannarale, *Astrophys. J.* **806**, 263 (2015), [arXiv:1405.7016 \[astro-ph.HE\]](https://arxiv.org/abs/1405.7016).
- [11] V. Tiwari and S. Fairhurst, *Astrophys. J. Lett.* **913**, L19 (2021), [arXiv:2011.04502 \[astro-ph.HE\]](https://arxiv.org/abs/2011.04502).
- [12] V. Tiwari, *Class. Quant. Grav.* **38**, 155007 (2021), [arXiv:2006.15047 \[astro-ph.HE\]](https://arxiv.org/abs/2006.15047).
- [13] D. Veske, I. Bartos, Z. Márka, and S. Márka, (2021), [arXiv:2105.13983 \[gr-qc\]](https://arxiv.org/abs/2105.13983).
- [14] S. Rinaldi and W. Del Pozzo, (2021), [arXiv:2109.05960 \[astro-ph.IM\]](https://arxiv.org/abs/2109.05960).
- [15] B. Edelman, Z. Doctor, J. Godfrey, and B. Farr, (2021), [arXiv:2109.06137 \[astro-ph.HE\]](https://arxiv.org/abs/2109.06137).
- [16] J. Powell, S. Stevenson, I. Mandel, and P. Tino, *Mon. Not. Roy. Astron. Soc.* **488**, 3810 (2019), [arXiv:1905.04825 \[astro-ph.HE\]](https://arxiv.org/abs/1905.04825).
- [17] B. Farr, D. E. Holz, and W. M. Farr, *Astrophys. J. Lett.* **854**, L9 (2018), [arXiv:1709.07896 \[astro-ph.HE\]](https://arxiv.org/abs/1709.07896).
- [18] T. Menne, “awkDE code,” <https://github.com/mennnthon/awkde>, with modifications at <https://github.com/JAMSADIQ/awkde> (2021), [Online; accessed Jul 2021].
- [19] I. Mandel, *Phys. Rev. D* **81**, 084029 (2010), [arXiv:0912.5531 \[astro-ph.HE\]](https://arxiv.org/abs/0912.5531).
- [20] D. Wysocki, J. Lange, and R. O’Shaughnessy, *Physical Review D* **100** (2019).
- [21] B. Efron, *The Annals of Statistics* **7**, 1 (1979).
- [22] C. Talbot and E. Thrane, *Astrophys. J.* **856**, 173 (2018), [arXiv:1801.02699 \[astro-ph.HE\]](https://arxiv.org/abs/1801.02699).
- [23] J. Sadiq and T. Dent, “PeakDetectionAlgorithm,” <https://github.com/JAMSADIQ/PeakDetectionAlgorithm> (2022), [Online; 2022].
- [24] R. Farmer, M. Renzo, S. E. de Mink, P. Marchant, and S. Justham, (2019), [10.3847/1538-4357/ab518b](https://doi.org/10.3847/1538-4357/ab518b), [arXiv:1910.12874 \[astro-ph.SR\]](https://arxiv.org/abs/1910.12874).
- [25] E. J. Baxter, D. Croon, S. D. McDermott, and J. Sakstein, *Astrophys. J. Lett.* **916**, L16 (2021), [arXiv:2104.02685 \[astro-ph.CO\]](https://arxiv.org/abs/2104.02685).
- [26] S. E. Woosley and A. Heger, *Astrophys. J. Lett.* **912**, L31 (2021), [arXiv:2103.07933 \[astro-ph.SR\]](https://arxiv.org/abs/2103.07933).
- [27] B. Silverman, *Density Estimation for Statistics and Data Analysis*, 1st ed. (Chapman and Hall, London, 1986).
- [28] G. R. Terrell and D. W. Scott, *The Annals of Statistics* **20**, 1236 (1992).
- [29] S. R. Sain and D. W. Scott, *Journal of the American Statistical Association* **91**, 1525 (1996).
- [30] B. Wang and X. Wang, “Bandwidth selection for weighted kernel density estimation,” (2011), [arXiv:0709.1616 \[stat.ME\]](https://arxiv.org/abs/0709.1616).
- [31] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, Springer Series in Statistics (Springer New York Inc., New York, NY, USA, 2001).
- [32] J. Veitch *et al.*, *Phys. Rev. D* **91**, 042003 (2015), [arXiv:1409.7215 \[gr-qc\]](https://arxiv.org/abs/1409.7215).
- [33] LVC, “GWTC-2 Data Release: Parameter Estimation Samples and Skymaps,” <https://dcc.ligo.org/LIGO-P2000223/public/> (2020), [Online; LIGO Document P2000223-v7].
- [34] D. Wysocki, “Calibrating semi-analytic VT’s to injections in O3a,” <https://dcc.ligo.org/LIGO-T2000432/public> (2020), [Online; LIGO Document T2000432-v1].
- [35] B. P. Abbott *et al.* (LIGO Scientific, Virgo), *Astrophys. J. Lett.* **882**, L24 (2019), [arXiv:1811.12940 \[astro-ph.HE\]](https://arxiv.org/abs/1811.12940).
- [36] C. Hoy and V. Raymond, “Bounded kdes,” (2021).
- [37] C. Hoy and V. Raymond, *SoftwareX* **15**, 100765 (2021), [arXiv:2006.06639 \[astro-ph.IM\]](https://arxiv.org/abs/2006.06639).
- [38] M. Fishbach and D. E. Holz, *Astrophys. J. Lett.* **851**, L25 (2017), [arXiv:1709.08584 \[astro-ph.HE\]](https://arxiv.org/abs/1709.08584).