Unmodelled clustering methods for gravitational wave populations of compact binary mergers

Jade Powell [1,2]★ Simon Stevenson [1,2] Ilya Mandel[2,3,4] and Peter Tiňo[5]

[1]*Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn, VIC 3122, Australia*
[2]*OzGrav: The ARC Centre of Excellence for Gravitational Wave Discovery, Melbourne VIC 3122, Australia*
[3]*School of Physics and Astronomy, Monash University, Clayton VIC 3800, Australia*
[4]*Institute of Gravitational Wave Astronomy and School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK*
[5]*School of Computer Science, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK*

ABSTRACT
The mass and spin distributions of compact binary gravitational-wave sources are currently uncertain due to complicated astrophysics involved in their formation. Multiple sub-populations of compact binaries representing different evolutionary scenarios may be present amongst sources detected by Advanced LIGO and Advanced Virgo. In addition to hierarchical modelling, unmodelled methods can aid in determining the number of sub-populations and their properties. In this paper, we apply Gaussian mixture model clustering to 1000 simulated gravitational-wave compact binary sources from a mixture of five sub-populations. Using both mass and spin as input parameters, we determine how many binary detections are needed to accurately determine the number of sub-populations and their mass and spin distributions. In the most difficult case that we consider, where two sub-populations have identical mass distributions but differ in their spin, which is poorly constrained by gravitational-wave detections, we find that ∼400 detections are needed before we can identify the correct number of sub-populations.

Key words: gravitational waves.

1 INTRODUCTION

The Advanced LIGO (aLIGO; The LIGO Scientific Collaboration et al. 2015) and Advanced Virgo (AdVirgo; Acernese et al. 2015) gravitational-wave detectors observed 10 stellar-mass binary black hole (BBH) mergers (Abbott et al. 2016; The LIGO Scientific Collaboration et al. 2018) and a binary neutron star (BNS) merger (Abbott et al. 2017) during the first two observing runs. The third observing run is currently underway, and several gravitational-wave detection candidates have been identified and circulated to electromagnetic partners via Gamma-ray Coordinates Network (GCN).[1] These merging compact binaries may have formed through several different formation mechanisms. Some of the possible formation mechanisms are classical isolated binary evolution (e.g. Belczynski et al. 2016), dynamical interactions in dense stellar environments (e.g. Heggie 1975), chemically homogeneous evolution (Mandel & de Mink 2016; Marchant et al. 2016), triple system formation (e.g. Antonini, Toonen & Hamers 2017; Rodriguez & Antonini 2018), or even a single event gravitationally lensed into multiple events (Broadhurst, Diego & Smoot

2019). Multiple compact binary gravitational-wave detections enable the study of the properties of populations of sources, which will constrain formation mechanisms (e.g. Stevenson, Ohme & Fairhurst 2015; Farr et al. 2017; Gerosa & Berti 2017; Stevenson, Berry & Mandel 2017; Vitale et al. 2017; Zevin et al. 2017; Talbot & Thrane 2018; Wysocki, Lange & O'Shaughnessy 2018a)

The rates, masses, and spins of BBHs can inform our understanding of their formation. If reliable models are available for different formation mechanisms, then hierarchical inference can be applied to compact objects detected in gravitational waves. This approach is optimal when models can be trusted, and allows the mixing ratios of different sub-populations and the properties of each population (e.g. the physics governing natal kicks received during core collapse or common envelope phases during binary evolution) to be measured (Farr et al. 2015).

Several studies have compared mass distributions to population synthesis models (Stevenson et al. 2015; Zevin et al. 2017; Gerosa et al. 2018; Taylor & Gerosa 2018; Wysocki et al. 2018b). For example, Barrett et al. (2018) argued that the parameters describing common envelope ejection efficiency, stellar wind strength, and natal kicks can all be measured to an accuracy of a few per cent with a thousand detections, provided the parametrized evolutionary model is accurate.

★ E-mail: dr.jade.powell@gmail.com
[1]https://gcn.gsfc.nasa.gov

The spin magnitudes and the alignment between the black hole spins and the orbital angular momentum carry information about the black hole binary formation mechanism. However, spin is difficult to constrain with gravitational-wave observations. Isotropic spin-orbit misalignment angles are expected for dynamically formed binaries (Heggie 1975; Bogdanović, Reynolds & Miller 2007; Rodriguez et al. 2016), but isolated binaries are expected to be preferentially aligned with the orbital angular momentum (e.g. Belczynski et al. 2016; Kushnir et al. 2016; Qin et al. 2019). To determine the distinguishability of these two formation mechanisms with gravitational-wave observations, hierarchical models have been applied to real and simulated gravitational-wave measurements of spin-orbit misalignment angles (Farr et al. 2017; Stevenson et al. 2017; Talbot & Thrane 2017; Vitale et al. 2017; Farr, Holz & Farr 2018; The LIGO Scientific Collaboration & The Virgo Collaboration 2018).

However, given the many modelling uncertainties, unmodelled or weakly modelled inference is a necessary back-up tool for studying source populations. Previous studies include fitting phenomenological population hyper-parameters to mass distributions, assuming the mass distribution is a power law inherited from the stellar initial mass function (Kovetz et al. 2017; The LIGO Scientific Collaboration & The Virgo Collaboration 2018; Wysocki et al. 2018a; Roulet & Zaldarriaga 2019), and including an upper mass gap and an excess of black holes near $40 \, M_{\odot}$ (Fishbach & Holz 2017; Talbot & Thrane 2018).

Unmodelled clustering techniques are particularly useful for interpreting population data when there is limited confidence in the available models. Mandel et al. (2015) argued that clustering can distinguish mock sub-populations of BNSs, neutron star black holes, and BBHs with tens of observations using only information about the masses of the two compact objects. Mandel et al. (2017) demonstrated that this is achievable for populations of compact binary mergers whose true mass distributions do not overlap (but whose measured properties do overlap because of measurement uncertainty). The method used in Mandel et al. (2017) involves reconstructing the observed mass distribution of merging compact binaries with a Gaussian process prior over a pixellated two-dimensional mass distribution with $\sim 10^2$ bins. Clusters are found in the reconstructed mass distribution with a 'water filling' algorithm (Mandel et al. 2017).

The clustering method of Mandel et al. (2017) scales poorly to a high-dimensional observable parameter space. The number of bins is exponential in the number of dimensions; $\sim 10^4$ bins will be needed in four dimensions. The 'water filling' clustering algorithm does not scale trivially to higher dimensions.

A Gaussian mixture model is an alternative clustering algorithm that scales well with the number of dimensions. A Gaussian mixture model fits a convex combination of multivariate Gaussian distributions to the input data. Wysocki (2017) apply a Gaussian mixture model to two simple examples. The first is a synthetic population of 30 binaries containing two sub-populations widely separated in the parameter space. One of the sub-populations has high mass and low spin, and the other has low mass and high spin. They predict the correct number of sub-populations but incorrectly predict the features of each sub-population when there are only 30 detections. The second example is 1000 perfect measurements from a single power-law distribution containing a mass gap.

Gaussian mixture models have also been applied to gravitational-wave detector noise transients, as finding different populations of noise transients can help identify their origin (Powell et al. 2015).

We expand on previous studies by applying clustering methods to sources sampled from a mixture of five complex sub-populations of simulated compact binaries. We model the imperfect inference on individual source parameters from their gravitational-wave signatures via realistic simulated posterior distributions. We apply a Gaussian mixture model (Pedregosa et al. 2011), assuming that the source population can be represented by a mixture of multivariate Gaussians. We apply this method simultaneously to both the mass and spin parameters for all of the binary systems considered. We consider how well our method can predict the correct properties of each sub-population. We show that the freedom in the choice of parametrization, the choice of cluster shape, and the choice of a distance metric can have a large effect on the results.

This paper is structured as follows. Section 2 describes the mock astrophysical sub-populations considered in this study. Section 3 explains how the mock inferred posteriors on the gravitational-wave parameters are produced. Section 4 presents the Gaussian mixture model method applied to the simulated sub-populations. The results are shown in Section 5, and a conclusion and discussion are given in Section 6.

2 SUB-POPULATIONS

In this study, we simulate five sub-populations of compact binaries. There are a total of 1000 simulated binaries with 20 per cent from each sub-population. The choice of 20 per cent is arbitrary, as we do not know what the true mixture fraction of sub-populations will be. This choice is not intended to be realistic; for examle, it does not respect the currently observed ratios between different binary types (The LIGO Scientific Collaboration et al. 2018). In Section 5, we show that varying the mixing ratios does not significantly affect our results.

We consider four parameters for each simulated binary. They are the companion masses m_1 and m_2, and the aligned spins χ_1 and χ_2, where

$$\chi_{1,2} = a_{1,2} \cos \theta_{1,2}. \tag{1}$$

The spin-orbit misalignment angle is given by θ, and a is the spin magnitude.

The true masses and spins of the compact objects in each sub-population are shown in Fig. 1, and further details of their ranges are given in Table 1. We consider our sub-populations as toy models that demonstrate our method.

The first three sub-populations are BNSs, neutron star black hole binaries (NSBHs), and BBHs (BBH1). The mass distributions are taken from Dominik et al. (2015) as in the previous studies of Mandel et al. (2015, 2017). The sub-populations were produced by the population synthesis code StarTrack (Belczynski et al. 2008) and down-selected to only include binaries in the aLIGO and AdVirgo detection range. Fig. 1 shows the clear gap in mass between the three sub-populations. We further assume that all compact objects in these three populations have zero spin.

The fourth sub-population (BBH2) has the same mass distribution as BBH1 but allows for the secondary companion to have aligned spin: a_1 is zero, and the secondary black hole spin a_2 is uniformly distributed between 0.7 and 0.95. For BBHs formed through common envelope evolution, the first born black hole is likely to have negligible spin (Qin et al. 2018; Bavera et al. 2019). This is because most of its angular momentum is stored in the envelope during the giant stage, and is removed during mass transfer or common envelope evolution. The spin of the second born black hole is determined by the effects of wind mass-loss and tides on
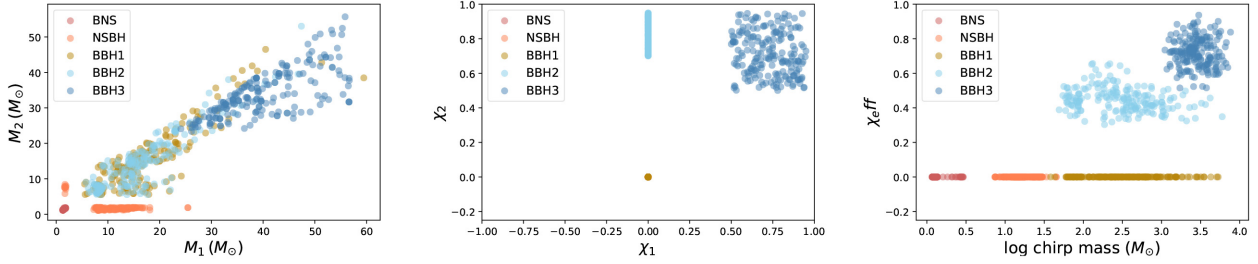
Figure 1. The true mass and spin values of 1000 simulated compact binaries. The different colours correspond to the different sub-populations. *Left*: The true values of the individual masses. Populations BBH1 and BBH2 have identical mass distributions. *Middle*: The true values of the aligned spins of the black hole sub-populations. The sub-populations containing neutron stars all have zero spin. *Right*: The true values converted to effective spin and log chirp mass.

Table 1. Details of the five sub-populations considered in this study. The first is non-spinning BNSs. The second is neutron star black hole binaries. The third and fourth sub-populations are black holes with the same mass distribution but different spin distributions. The fifth population contains higher mass black holes with larger spins.

Short name	Source type	Mass distribution	m_1 90% range	m_2 90% range	Spin distribution
BNS	Neutron star binaries	Dominik et al. (2015)	1.3–1.32 $M_\odot$	1.2–1.3 $M_\odot$	$\chi_1 = \chi_2 = 0$
NSBH	NS–BH binaries	Dominik et al. (2015)	7–15 $M_\odot$	1.3–1.9 $M_\odot$	$\chi_1 = \chi_2 = 0$
BBH1	Black hole binaries	Dominik et al. (2015)	7–28 $M_\odot$	5–24 $M_\odot$	$\chi_1 = \chi_2 = 0$
BBH2	Black hole binaries	Dominik et al. (2015)	7–26 $M_\odot$	5–23 $M_\odot$	$\chi_1 = 0, \chi_2 = U(0.7, 0.95)$
BBH3	Black hole binaries	Mandel & de Mink (2016)	29–56 $M_\odot$	25–44 $M_\odot$	$\chi_1 = \chi_2 = U(0.5, 1.0)$

its progenitor, a helium star (Kushnir et al. 2016; Gerosa et al. 2018; Qin et al. 2018; Zaldarriaga, Kushnir & Kollmeier 2018; Bavera et al. 2019). If the binary has a sufficiently short orbital period, the helium star may be spun up through tides and form a rapidly spinning black hole. Otherwise, in binaries with larger orbital periods, the helium star is not spun up through tides, and loses angular momentum through stellar winds. We use this to motivate the spin values of BBH1 and BBH2. Distinguishing between BBH1 and BBH2 is expected to be the most difficult example considered in this study due to the poor constraints on gravitational-wave spin measurements.

BBHs formed through binary evolution are expected to have their spins preferentially aligned with the orbital angular momentum. The spins are perfectly aligned in both sub-populations that have spins, i.e. $\theta = 0$ for all our binary systems. Misalignments are fairly uncertain, due to uncertainties both in the magnitude of black hole kicks, and in realignment processes through binary evolution (Gerosa et al. 2018; Wysocki et al. 2018b). We do not expect the assumption of spin-orbit alignment to have a significant effect on our results as the distributions would still be reasonably well separated if a more moderate amount of misalignment was included, as expected in more realistic models.

The fifth sub-population (BBH3) consists of black holes with a higher mass distribution based on the results of Mandel & de Mink (2016), where they consider merging BBHs formed through chemically homogeneous evolution in short-period stellar binaries. The individual masses are distributed between approximately 25 and 50 $M_\odot$. The aligned spins of the individual black holes are motivated by the results of Marchant et al. (2016) and are uniformly distributed between $\chi = 0.5$ and $\chi = 1.0$.

3 POSTERIORS

Bayesian parameter estimation is applied to data containing detected gravitational-wave signals to produce posterior distributions for the astrophysical parameters of the source (Veitch et al. 2015).

For compact binary signals, the parameters include the mass, spin, eccentricity, distance, inclination, and sky position. Bayesian inference is computationally expensive for a large number of detections. Therefore, for this study, we produce realistic artificial posteriors for mass parameters as in Mandel et al. (2017), and spin parameters as in Stevenson et al. (2017).

For the mass posteriors, we first sample values for m_1 and m_2 from each of the five sub-populations described in Section 2. For each binary system, we then draw a signal-to-noise ratio (SNR) ρ value from a $p(\rho) \propto \rho^{-4}$ distribution (Schutz 2011), where a network SNR $\rho \geq 12$ is used as a threshold for detectability. The mass parameters are converted into a chirp mass M_c given by

$$M_c = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}, \qquad (2)$$

and a symmetric mass ratio η given by

$$\eta = \frac{m_1 m_2}{(m_1 + m_2)^2}. \qquad (3)$$

We then calculate the posteriors for the masses using the method in Mandel et al. (2017). The posteriors are generated in the chirp mass parameter space as

$$M_c = M_c^T \left[1 + \alpha \frac{12}{\rho}(r_0 + r) \right] \qquad (4)$$

and the symmetric mass ratio as

$$\eta = \eta^T \left[1 + 0.03 \frac{12}{\rho}(r_0 + r) \right]. \qquad (5)$$

Here, M_c^T and η^T are the true values, r_0 is a random number drawn from a standard normal distribution that represents the shift of the mean of the posterior with respect to the true value, and r is an array of random numbers from the standard normal distribution. As in Mandel et al. (2017), the parameter α, which determines the width of the posterior distribution, is motivated by the previous studies of Littenberg et al. (2015) and Mandel et al. (2015), and has values of 0.01, 0.03, and 0.1 when $\eta^T > 0.1$, $0.1 > \eta^T > 0.05$, and $0.05 > \eta^T$,

respectively. The posterior samples on M_c and η are then converted to samples back in the m_1 and m_2 parameter space.

We then generate the spin posteriors using the method of Stevenson et al. (2017). As for the mass posteriors, we sample values for χ_1 and χ_2 from the relevant distributions from each of the five sub-populations. Those values can then be used to determine the effective spin $\chi_{\rm eff}$ (Damour 2001; Racine 2008; Ajith et al. 2011) defined as

$$\chi_{\rm eff} = \frac{a_1 \cos\theta_1 + q a_2 \cos\theta_2}{(1+q)}, \tag{6}$$

where the mass ratio $q = m_2/m_1$. The posteriors are generated in the $\chi_{\rm eff}$ and χ_1 parameter space as

$$\chi_{\rm eff} = \chi_{\rm eff}^T + \left[\beta \frac{12}{\rho}(r_0 + r)\right], \tag{7}$$

$$\chi_1 = \chi_1^T + \left[\gamma \frac{12}{\rho}(r_0 + r)\right], \tag{8}$$

where $\beta = 0.1$ and $\gamma = 0.2$, as in Stevenson et al. (2017), and limits of -1 and 1 are applied to both parameters. A posterior for the χ_2 parameter is then determined from the posteriors for $\chi_{\rm eff}$, χ_1, and q. The same number of posterior samples, between 200 and 1000, are used for the spin parameters of a given event.

The 90 per cent confidence interval contours for five representative posterior distributions from each of the three BBH sub-populations are shown in Fig. 2. The mass posteriors show the typically expected banana shape. Their size is related to the SNR of the signal. The aligned spin posteriors are much wider than the mass posteriors, as spin is poorly constrained by gravitational-wave detections.

In a generic BBH, the black hole spins will not be perfectly aligned with the orbital angular momentum. In these systems, the black hole spins will precess (Apostolatos et al. 1994; Gerosa et al. 2014). This causes the distribution of spin-orbit misalignment angles $\cos\theta_{1,2}$ to vary with orbital frequency (or equivalently, gravitational-wave frequency). The distribution of $\chi_{\rm eff}$ is approximately constant at the second post-Newtonian order (Racine 2008). An isotropic distribution of spins is expected to remain isotropic through post-Newtonian evolution (Schnittman 2004; Bogdanović et al. 2007). These considerations will be important in choosing the parametrization for clustering on real gravitational-wave observations.

4 CLUSTERING METHOD

In this study, we consider a Gaussian mixture model as implemented in Scikit learn (Pedregosa et al. 2011). A Gaussian mixture model fits a linear combination of multivariate Gaussian distributions to the input data. Our implementation of this method requires smooth inputs for clustering, so it is not possible to use all of the posterior samples as input. We therefore represent each observation by a set of estimators: the median and 90 per cent confidence limit values of each marginalized one-dimensional posterior. Thus, the input data x for an observation with n inferred parameters is a vector of length $3 \times n$. An n-dimensional Gaussian probability density is given by

$$p(x|\mu, \Sigma) = \frac{1}{\sqrt{(2\pi)^n \det\Sigma}} \exp\left[-\frac{1}{2}(x-\mu)^T \Sigma^{-1}(x-\mu)\right], \tag{9}$$

where μ is the mean, and Σ is the covariance matrix. The likelihood function for a single observation under a K-component Gaussian

mixture model is then given by

$$\text{GMM}(x|w, \{\mu, \Sigma\}) = \sum_{k=1}^{K} w_k\, p(x|\mu_k, \Sigma_k), \tag{10}$$

where w_k are the mixture weights equal to marginal probabilities of mixture components. The free parameters for each Gaussian are its weight (up to an overall normalization), mean, and covariance matrix. Therefore, we need to find the optimum parameters that maximize the likelihood in order to fit the mixture of Gaussians to the observed compact binary population. This likelihood for the full set of N observations $X = \{X_1...X_N\}$ is given by

$$p(X|w, \mu, \Sigma) = \prod_{j=1}^{N} \text{GMM}(X_j|w, \mu, \Sigma). \tag{11}$$

An expectation maximization technique (Dempster, Laird & Rubin 1977) is then used to find a maximum likelihood estimate that determines the correct values of the means, weights, and covariance matrices for a given number of Gaussians. In this method, the weights, means, and covariances are first randomly initialized. Then the expectation and maximization steps are iterated repeatedly so that the likelihood of the data increases at the end of each step.

To determine the correct number of Gaussians, we minimize the Bayesian Information Criterion (BIC) given by

$$\text{BIC} = -2\ln\mathcal{L} + k\ln N, \tag{12}$$

where k is the number of free parameters to be estimated and $\mathcal{L}$ is the maximized value of the likelihood of the best-fitting model. BIC adds a penalty to models with larger numbers of parameters to avoid over fitting. The maximum number of Gaussians we consider for this method is 10. This method can produce results in a few seconds for very large numbers of detected compact binaries. Other mixture model methods that can be used for model selection include infinite mixtures and bottom-up growing and top-down pruning techniques (Rasmussen 2000; Figueiredo & Jain 2002; Del Pozzo et al. 2018).

5 RESULTS

Here, we report on the results of applying a Gaussian mixture model to our mock population. We are particularly interested in the number of clusters re-constructed (does it equal the five modelled subpopulations) and the weights assigned to each cluster (does it equal 20 per cent as in the model). We also track the fraction of events from the same sub-population that are clustered together. These results are given in Table 2.

We vary the choices of parameters used for clustering. Clustering requires a choice of a metric on the parameter space in order to define a distance between clusters and a choice for the shape of each cluster. Neither can be determined from first principles. Whilst the assumed Gaussian probability density with covariance matrices define the cluster shape and a (Euclidean) distance metric, changing the parametrization effectively changes the cluster shape and distance; for example, transforming to logarithmic coordinates is equivalent to assuming a lognormal cluster shape rather than a normal one. We thus consider using log chirp mass and the effective spin as clustering parameters, as these are better determined from observations than individual masses and spins. We also include results using only mass, spin, and the BBH sub-populations. The results are given in Table 2.

Although optimality is difficult to define, we find that for our particular population of events, using individual masses and spins
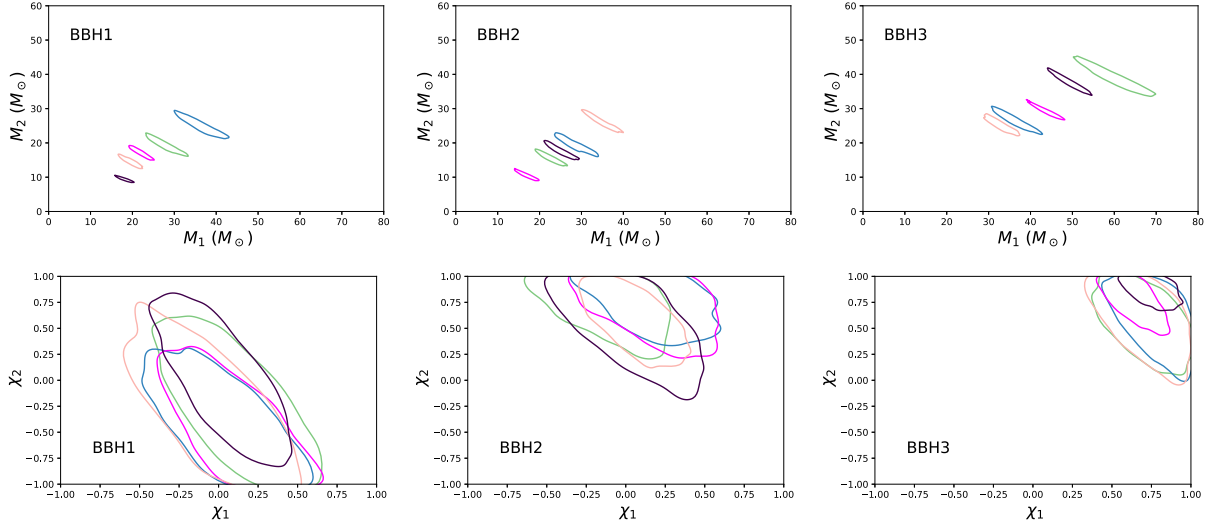
Figure 2. Five representative 90 per cent confidence interval contours for the mass and spin posteriors of the three different BBH populations. The mass posteriors show the typical banana shape and are better constrained than the spin. BBH1 and BBH2 have the same mass distributions but differ in their aligned spins.

Table 2. Gaussian mixture model results for different combinations of parameters. The best results are obtained when parametrizing the full population with m_1, m_2, χ_1, and χ_2.

Parameters	No. predicted sub-populations	BNS	NSBH	BBH1	BBH2	BBH3
m_1, m_2	4	98% in the same cluster	99% in the same cluster	Mixed with BBH2	Mixed with BBH 1	79% in the same cluster
χ_1, χ_2	4	Mixed with BBH1	95% in the same cluster	Mixed with BNS	78% in the same cluster	57% in the same cluster
m_1, m_2, χ_1, χ_2	5	99% in the same cluster	100% in the same cluster	80% in the same cluster	68% in the same cluster	81% in the same cluster
$\log M_c$, χ_{eff}	5	92% in the same cluster	92% in the same cluster	Mixed with BBH2	Mixed with BBH1	91% in the same cluster
BBH only, m_1, m_2, χ_1, χ_2	3			74% in the same cluster	47% in the same cluster	68% in the same cluster
BBH only, $\log M_c$, χ_{eff}	3			Mixed with BBH2	Mixed with BBH1	92% in the same cluster

as clustering parameters yields robust results. The number of sub-populations is generally correctly determined for all parameter options. As expected, only four sub-populations are found when using the mass parameters alone, since BBH1 and BBH2 have identical mass distributions. When using spins alone, similar results to the combined spin and mass case are found for sub-population BBH2, however only around half of the BBH3 binaries are in the same sub-population, with the other half mixed in with the other sub-populations. Two of the sub-populations with zero spin, BNS and BBH1, were mixed as expected. However, the NSBH sub-population that also has zero spin was still distinguishable from the other zero spin sub-populations. This is due to the mass ratio value of the NSBH sub-population changing the shape of the posteriors in the individual spin parameter space. Log chirp mass and effective spin also lead to BBH1 and BBH2 being mixed into two classes that are split into lower and higher masses, rather than by spins, despite their different spin distributions. This may be partly related to a partial artificial breaking of the spin-degeneracy in our spin measurement uncertainty model.

When using only the binaries in the three BBH sub-populations, the results show that there is a larger mixing between sub-

populations than when the BNS and NSBH sub-populations are included. Using all possible compact binaries and the full set of parameters m_1, m_2, χ_1, and χ_2 yields the most information, and the Gaussian mixture model performs better with a larger set of input data. Now that the best parametrization has been identified, the rest of the results in this section use only χ_1, χ_2, m_1, and m_2 as input parameters.

We investigate how many detections are needed before the correct number of sub-populations can be identified. We detect the binary signals in a random order starting with 30 detections, increased in steps of five detections, and apply the Gaussian mixture model after each step. The results are shown in Fig. 3. With only 30 detections, the smallest number we consider, we can distinguish between the BNS, NSBH, and BBH sub-populations. The three BBH sub-populations are then split into a large number of Gaussians by the mass of the binaries. This is because the masses are better constrained by gravitational-wave signatures than the spin; therefore, a larger number of spin posteriors are required before the spin has a significant influence on the results. Accidental apparent clustering of a small number of observations in a multidimensional space leads to an overestimate of the number of clusters for $\lesssim 100$
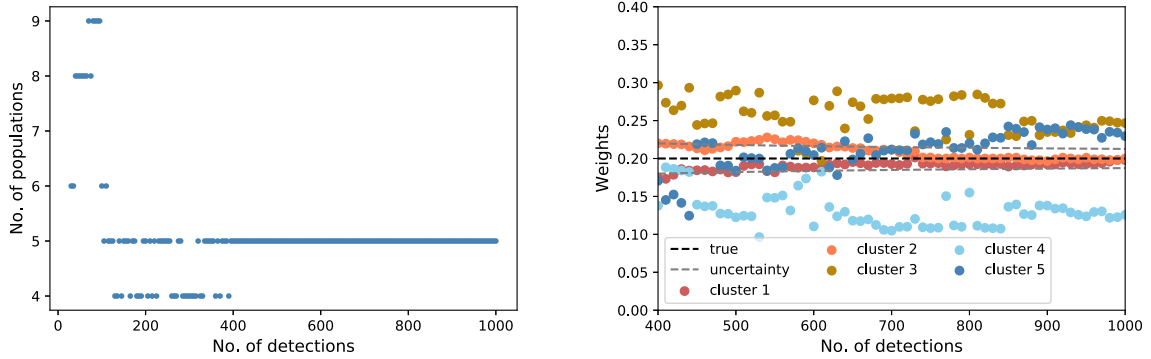
Figure 3. *Left*: The number of sub-populations estimated by the Gaussian mixture model parametrized with χ_1, χ_2, m_1, and m_2 as the number of detections is increased. Approximately 400 detections are needed to estimate the correct number of sub-populations. *Right*: The estimated weights of the five clusters after the correct number of clusters is determined.

detections. After 105 detections, the binaries in the BBH3 sub-population are correctly grouped together by the Gaussian mixture model, but the BBH1 and BBH2 sub-populations are still mixed either into the same sub-population or into two sub-populations that are split by their mass. After 400 detections, we can start to distinguish between the more difficult case of the BBH1 and BBH2 sub-populations.

In Fig. 3, we also show the weight values after the correct number of sub-populations has been determined. The correct number for all sub-populations should be 0.2 with a multinomial counting uncertainty that is given by

$$\text{uncertainty} = \sqrt{\frac{N_{\text{pop}} - 1}{N_{\text{pop}}^2 N_{\text{det}}}} = 0.4 N_{\text{det}}^{-0.5}, \tag{13}$$

where $N_{\text{pop}} = 5$ is the number of sub-populations, and N_{det} is the number of detections. The BNS and NSBH sub-populations have the correct weight values. There is a larger error in the weights of the three BBH sub-populations due to some mixing between the different sub-populations. This occurs due to the poor gravitational-wave spin measurements and two of the sub-populations only differing in their spin values.

After finding the correct number of sub-populations, we want to know the mass and spin distributions of each individual sub-population, as this will aid investigations into differences in their formation mechanisms. The individual mass and spin distributions determined by this method after 400 detections are shown in Fig. 4.

In cluster 1, associated with the BNS sub-population, both masses have distributions expected for neutron stars and both of the aligned spin distributions are centred on zero. In cluster 2, associated with the NSBH sub-population, mass distributions show that m_1 is typically a low-mass black hole and that m_2 is a neutron star. The spin distributions of cluster 2 (NSBH) and cluster 3 (BBH1) are centred on zero, as expected. Cluster 5, associated with the BBH3 sub-population, contains systems with high masses and high spins, whilst cluster 4 (BBH2) exhibits a χ_2 distribution favouring high spin whilst the χ_1 distribution is centred on zero.

We investigate the impact of different mixing ratios in the mock population by considering two other population models consisting of the same five sub-populations present with different relative mixing ratios. In all cases, we evaluate the clustering results after 620 detections. The first new mixture, which we refer to as mixture 2, consists of a much smaller number of BNS and NSBH signals as expected from current gravitational-wave detections. The total

population in this mixture consists of 6.5 per cent BNS, 6.5 per cent NSBH, 29 per cent BBH1, 29 per cent BBH2, and 29 per cent BBH3 signals. The second new mixture, which we refer to as mixture 3, has a smaller number of BNS, NSBH, and BBH in one of the BBH sub-populations. This mixture consists of 16 per cent BNS, 16 per cent NSBH, 16 per cent BBH3, 26 per cent BBH1, and 26 per cent BBH2 signals.

The results for the different mixtures are shown in Table 3. Our clustering algorithm determines the correct number of sub-populations for all of the mixing ratios considered. For mixture 2, where there is now a much smaller number of BNS and NSBH detections, the increase in the ratio of detections from all three BBH sub-populations leads to an improved clustering result for those three sub-populations. For mixture 3, where we also decrease the BBH3 population, the accuracy of the results is similar to the equal mixture case.

6 CONCLUSIONS

Measuring the number, properties, and shapes of the sub-populations associated with different source types and formation channels of merging compact binaries detected through their gravitational-wave signatures will assist in the astrophysical interpretation of these observations. We explore clustering on simulated distributions for the masses and spins of binary systems from a mixture of five sub-populations, using mock inference on their individual parameters from gravitational-wave observations. We demonstrate that Gaussian mixture model clustering on the full set of observations in the four-dimensional space of source parameters makes it possible to distinguish the multiple sub-populations present after 400 detections and determine the mass and spin distributions of each sub-population.

Gaussian mixture model clustering is robust and computationally efficient. It scales well to higher-dimensional parameter spaces, but it requires a partial loss of the posterior information, as only a few estimators rather than the full event posteriors are used for clustering. On the other hand, as with any clustering technique, it relies on the choice of a distance metric in the parameter space. It also assumes a specific parametrized shape for the clusters in the mixture model. Although the method performs robustly on the mock populations we simulate, which are not multivariate Gaussians, it may be suboptimal for some non-Gaussian distribution shapes. This method requires an upper limit on the number of sub-populations; we use the BIC to determine the optimal number of sub-populations.
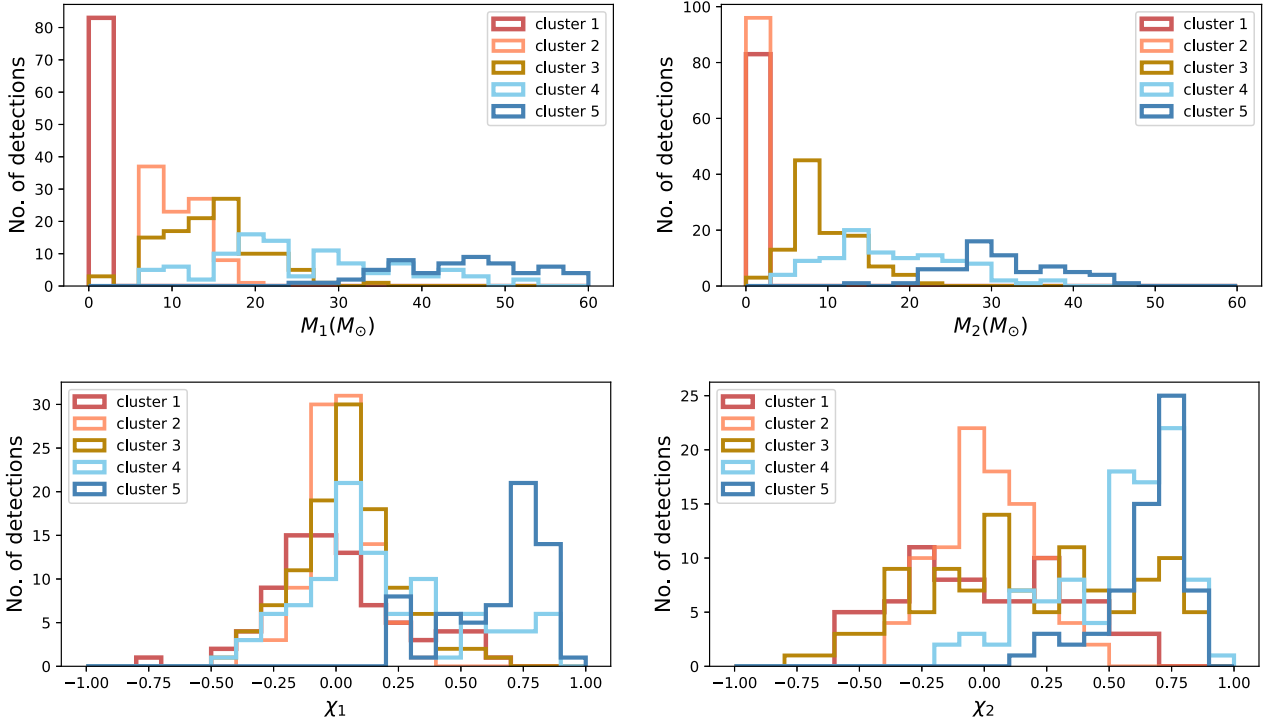
Figure 4. The individual mass and spin distributions (binned posterior medians) of the sub-populations identified by the Gaussian mixture model after 400 detections. *Top left*: The five m_1 distributions. *Top right*: The m_2 distributions. *Bottom left*: The χ_1 distributions. Four of the distributions are clearly centred on zero, whilst one distribution has higher values. *Bottom right*: The χ_2 distributions. Two of the distributions have higher χ_2 values than the other three, which are centred on zero.

Table 3. The results after 620 detections when analysing populations consisting of a mixture of the same five sub-populations with different mixing ratios. Mixture 2 contains a much smaller number of BNS and NSBH. Mixture 3 has fewer detections in the BBH3 sub-population as well BNS and NSBH. The Gaussian mixture model isolates five sub-populations for all mixtures considered.

Mixtures	No. predicted	BNS	NSBH	BBH1	BBH2	BBH3
Equal ratio	5	99% in the same cluster	100% in the same cluster	86% in the same cluster	69% in the same cluster	64% in the same cluster
Mixture 2	5	100% in the same cluster	100% in the same cluster	81% in the same cluster	79% in the same cluster	91% in the same cluster
Mixture 3	5	99% in the same cluster	100% in the same cluster	79% in the same cluster	79% in the same cluster	65% in the same cluster

Finally, this method does not allow us to produce uncertainty estimates on the number of sub-populations and their weights; however, this is a general challenge for any clustering method that relies on an arbitrary distance metric, and is not a specific flaw of the current technique.

The difficulty in distinguishing sub-populations depends on the similarity of the parameters of the binaries in the different sub-populations, and the exact number of observations necessary for this therefore depends on the population properties. The most difficult case we consider here is two sub-populations of BBHs that have identical mass distributions and differ only in one of their spin components. We do not incorporate selection effects in our analysis (Farr et al. 2015; Mandel, Farr & Gair 2019). We also do not consider the effect of lower significance gravitational-wave detections, which may be produced by transient noise that could not be distinguished from gravitational-wave sources (Gaebel et al. 2019).

REFERENCES

Abbott B. P. et al., 2016, PhRvX, 6, 041015
Abbott B. P. et al., 2017, Phys. Rev. Lett., 119, 161101
Acernese F. et al., 2015, Class. Quantum Gravity, 32, 024001
Ajith P. et al., 2011, Phys. Rev. Lett., 106, 241101
Antonini F., Toonen S., Hamers A. S., 2017, ApJ, 841, 77
Apostolatos T. A., Cutler C., Sussman G. J., Thorne K. S., 1994, Phys. Rev. D, 49, 6274

Barrett J. W., Gaebel S. M., Neijssel C. J., Vigna-Gómez A., Stevenson S., Berry C. P. L., Farr W. M., Mandel I., 2018, MNRAS, 477, 4685

Bavera S. S. et al., 2019, preprint (arXiv:1906.12257)

Belczynski K., Kalogera V., Rasio F. A., Taam R. E., Zezas A., Bulik T., Maccarone T. J., Ivanova N., 2008, ApJS, 174, 223

Belczynski K., Holz D. E., Bulik T., O'Shaughnessy R., 2016, Nature, 534, 512

Bogdanović T., Reynolds C. S., Miller M. C., 2007, ApJ, 661, L147

Broadhurst T., Diego J. M., Smoot George F. I., 2019, preprint (arXiv:1901.03190)

Damour T., 2001, Phys. Rev. D, 64, 124013

Del Pozzo W., Berry C. P. L., Ghosh A., Haines T. S. F., Singer L. P., Vecchio A., 2018, MNRAS, 479, 601

Dempster A. P., Laird N. M., Rubin D. B., 1977, J. R. Stat. Soc. Ser. B (Methodol.), 39, 1

Dominik M. et al., 2015, ApJ, 806, 263

Farr B., Holz D. E., Farr W. M., 2018, ApJ, 854, L9

Farr W. M., Gair J. R., Mandel I., Cutler C., 2015, Phys. Rev. D, 91, 023005

Farr W. M., Stevenson S., Miller M. C., Mandel I., Farr B., Vecchio A., 2017, Nature, 548, 426

Figueiredo M. A. T., Jain A. K., 2002, IEEE Trans. Pattern Anal. Mach. Intell., 24, 381

Fishbach M., Holz D. E., 2017, ApJ, 851, L25

Gaebel S. M., Veitch J., Dent T., Farr W. M., 2019, MNRAS, 484, 4008

Gerosa D., Berti E., 2017, Phys. Rev. D, 95, 124046

Gerosa D., O'Shaughnessy R., Kesden M., Berti E., Sperhake U., 2014, Phys. Rev. D, 89, 124025

Gerosa D., Berti E., O'Shaughnessy R., Belczynski K., Kesden M., Wysocki D., Gladysz W., 2018, Phys. Rev. D, 98, 084036

Heggie D. C., 1975, MNRAS, 173, 729

Kovetz E. D., Cholis I., Breysse P. C., Kamionkowski M., 2017, Phys. Rev. D, 95, 103010

Kushnir D., Zaldarriaga M., Kollmeier J. A., Waldman R., 2016, MNRAS, 462, 844

Littenberg T. B., Farr B., Coughlin S., Kalogera V., Holz D. E., 2015, ApJ, 807, L24

Mandel I., de Mink S. E., 2016, MNRAS, 458, 2634

Mandel I., Haster C.-J., Dominik M., Belczynski K., 2015, MNRAS, 450, L85

Mandel I., Farr W. M., Colonna A., Stevenson S., Tiňo P., Veitch J., 2017, MNRAS, 465, 3254

Mandel I., Farr W. M., Gair J. R., 2019, MNRAS, 486, 1086

Marchant P., Langer N., Podsiadlowski P., Tauris T. M., Moriya T. J., 2016, A&A, 588, A50

Pedregosa F. et al., 2011, J. Mach. Learn. Res., 12, 2825

Powell J., Trifirò D., Cuoco E., Heng I. S., Cavaglià M., 2015, Class. Quantum Gravity, 32, 215012

Qin Y., Fragos T., Meynet G., Andrews J., Sørensen M., Song H. F., 2018, A&A, 616, A28

Qin Y., Marchant P., Fragos T., Meynet G., Kalogera V., 2019, ApJ, 870, L18

Racine E., 2008, Phys. Rev. D, 78, 044021

Rasmussen C. E., 2000, in Solla S. A., Leen T. K., Müller K., eds, Advances in Neural Information Processing Systems 12. MIT Press, Cambridge, p. 554

Rodriguez C. L., Antonini F., 2018, ApJ, 863, 7

Rodriguez C. L., Zevin M., Pankow C., Kalogera V., Rasio F. A., 2016, ApJ, 832, L2

Roulet J., Zaldarriaga M., 2019, MNRAS, 484, 4216

Schnittman J. D., 2004, Phys. Rev. D, 70, 124020

Schutz B. F., 2011, Class. Quantum Gravity, 28, 125023

Stevenson S., Ohme F., Fairhurst S., 2015, ApJ, 810, 58

Stevenson S., Berry C. P. L., Mandel I., 2017, MNRAS, 471, 2801

Talbot C., Thrane E., 2017, Phys. Rev. D, 96, 023012

Talbot C., Thrane E., 2018, ApJ, 856, 173

Taylor S. R., Gerosa D., 2018, Phys. Rev. D, 98, 083017

The LIGO Scientific Collaboration et al., 2018, preprint (arXiv:1811.12907)

The LIGO Scientific Collaboration, The Virgo Collaboration, 2018, preprint (arXiv:1811.12940)

The LIGO Scientific Collaboration et al., 2015, Class. Quantum Gravity, 32, 074001

Veitch J. et al., 2015, Phys. Rev. D, 91, 042003

Vitale S., Lynch R., Sturani R., Graff P., 2017, Class. Quantum Gravity, 34, 03LT01

Wysocki D., 2017, preprint (arXiv:1712.02643)

Wysocki D., Lange J., O'Shaughnessy R., 2018a, preprint (arXiv:1805.06442)

Wysocki D., Gerosa D., O'Shaughnessy R., Belczynski K., Gladysz W., Berti E., Kesden M., Holz D. E., 2018b, Phys. Rev. D, 97, 043014

Zaldarriaga M., Kushnir D., Kollmeier J. A., 2018, MNRAS, 473, 4174

Zevin M., Pankow C., Rodriguez C. L., Sampson L., Chase E., Kalogera V., Rasio F. A., 2017, ApJ, 846, 82

This paper has been typeset from a TeX/LaTeX file prepared by the author.