

Model-independent inference on compact-binary observations

Mandel, Ilya; Farr, Will M.; Colonna, Andrea; Stevenson, Simon; Tiño, Peter; Veitch, John

DOI:

[10.1093/mnras/stw2883](https://doi.org/10.1093/mnras/stw2883)

License:

Other (please specify with Rights Statement)

Document Version

Publisher's PDF, also known as Version of record

Citation for published version (Harvard):

Mandel, I, Farr, WM, Colonna, A, Stevenson, S, Tiño, P & Veitch, J 2017, 'Model-independent inference on compact-binary observations', *Royal Astronomical Society. Monthly Notices*, vol. 465, no. 3, pp. 3254-3260. <https://doi.org/10.1093/mnras/stw2883>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

This article has been accepted for publication in *Monthly Notices of the Royal Astronomical Society* ©: 2016 The Authors Published by Oxford University Press on behalf of the Royal Astronomical Society. All rights reserved.

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Model-independent inference on compact-binary observations

Ilya Mandel,¹★ Will M. Farr,¹ Andrea Colonna,² Simon Stevenson,¹ Peter Tiño²
and John Veitch¹

¹*School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK*

²*School of Computer Science, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK*

Accepted 2016 November 4. Received 2016 October 27; in original form 2016 August 28

ABSTRACT

The recent advanced LIGO detections of gravitational waves from merging binary black holes enhance the prospect of exploring binary evolution via gravitational-wave observations of a population of compact-object binaries. In the face of uncertainty about binary formation models, model-independent inference provides an appealing alternative to comparisons between observed and modelled populations. We describe a procedure for clustering in the multidimensional parameter space of observations that are subject to significant measurement errors. We apply this procedure to a mock data set of population-synthesis predictions for the masses of merging compact binaries convolved with realistic measurement uncertainties, and demonstrate that we can accurately distinguish subpopulations of binary neutron stars, binary black holes, and mixed neutron star–black hole binaries with tens of observations.

Key words: gravitational waves – binaries: close – stars: black holes – stars: neutron.

1 INTRODUCTION

The advanced LIGO detectors (Aasi et al. 2015) observed the first gravitational waves from a merger of two black holes (BHs), GW150914, on 14 September 2015 (Abbott et al. 2016b). This discovery was followed in a few months by another BH–BH merger detection, GW151226 (Abbott et al. 2016d), and a further likely BH–BH candidate, LVT151012 (Abbott et al. 2016a). The BH–BH merger rate inferred from these events implies that tens to hundreds of detections are likely over the next few years (Abbott et al. 2016a,b). Meanwhile, both massive binary evolution models and observations of Galactic binary pulsars and short gamma-ray bursts suggest that gravitational-wave detections of mergers of two neutron stars (NSs) and mergers of mixed NS–BH binaries are also likely in the coming years (see Abadie et al. 2010, for a review).

Multiple observations should make it possible to address the inverse problem of gravitational-wave astrophysics: to study the currently uncertain massive stellar binary evolution through its evolution end products — the population of merging compact remnants. One approach to this problem involves creating forward models of binary evolution, e.g. via population synthesis Monte Carlo simulations (see Postnov & Yungelson 2014, for a review), and comparing them with the observed population to constrain the input assumptions, such as the common-envelope physics (Ivanova et al. 2013). This approach has been advocated by Bulik & Belczyński (2003), Mandel & O’Shaughnessy (2010), O’Shaughnessy (2013), and Stevenson, Ohme & Fairhurst (2015), among others.

While this approach is very promising, existing binary evolution models may not correctly encapsulate the full range of physical uncertainties (e.g. Dominik et al. 2012; Mennekens & Vanbeveren 2014; Belczynski et al. 2016; Eldridge & Stanway 2016; Lipunov et al. 2016). Moreover, some of the merging compact binaries could form through channels other than isolated binary evolution via the common-envelope phase, including chemically homogeneous evolution in very close binaries (Mandel & de Mink 2016; Marchant et al. 2016; de Mink & Mandel 2016), dynamical formation in globular clusters, young stellar clusters, or galactic nuclei (Rodríguez et al. 2016; Mapelli 2016; Bartos et al. 2016; Stone, Metzger & Haiman 2016), mergers of Population III remnants (Inayoshi et al. 2016) or even primordial BH mergers (Bird et al. 2016). In the possible presence of both systematic model uncertainty and confusion from different formation channels, a model-independent approach to learning from the observed population is desirable.

Mandel et al. (2015) proposed that clustering on the parameters of observed merging compact-object binaries could provide useful model-independent information about the population. This clustering is greatly complicated by the limited accuracy with which the masses and spins of merging binaries can be inferred from gravitational-wave observations (e.g. Littenberg et al. 2015; Veitch et al. 2015a; Abbott et al. 2016c). Nevertheless, Mandel et al. (2015) suggested that for astrophysically plausible binary populations and realistic measurement uncertainties, a few tens to a few hundred detections should be sufficient to cluster merging binaries into NS–NS, BH–BH, and NS–BH subpopulations, estimating their relative rates to within Poisson uncertainty.

★ E-mail: imandel@star.sr.bham.ac.uk




Figure 1. Masses of merging compact-object binaries as simulated via population synthesis. The lower left-hand corner is occupied by NS–NS merging binaries (yellow), the upper right-hand corner by the more massive BH–BH systems (blue), while the NS–BH population (orange) is asymmetric.

This paper describes specific algorithms for clustering on the observed merging compact-binary population in the presence of significant measurement uncertainty. We show how the clustering could proceed in practice when subpopulations with distinct mass parameters are brought into contact once the underlying mass distributions are convolved with measurement errors. We demonstrate the accuracy of the analytical predictions of Mandel et al. (2015) with a quantitative study. Our approach can be trivially extended to include other parameters such as spin magnitudes and spin tilt angles.

2 BINARY POPULATION

We analyse a realistic population of compact-object binaries produced with a population synthesis code that evolves binaries from zero-age main-sequence stars through stellar evolution, mass transfer including a possible common-envelope phase, wind-driven mass-loss, supernovae, and eventual gravitational-wave-driven mergers. For the ease of comparison, we use the same simulated binary data set as in Mandel et al. (2015). This data set was constructed with the `STARTRACK` code (Belczynski et al. 2008), using the ‘Standard’ model B of Dominik et al. (2012), including the rapid supernova engine (Belczynski et al. 2012; Fryer et al. 2012), down-selected to binaries potentially detectable by the advanced-detector network, as estimated by Dominik et al. (2015). A number of parameters governing binary evolution are highly uncertain, including wind-driven and luminous blue variable mass-loss rates (e.g. Vink, de Koter & Lamers 2001; Mennekens & Vanbeveren 2014), mass transfer efficiency (e.g. de Mink, Pols & Hilditch 2007), common-envelope physics (e.g. Ivanova et al. 2013), and BH natal kicks (e.g. Repetto & Nelemans 2015; Mandel 2016). Therefore, this model should be viewed only as a realistic illustration for the model-independent inference technique.

The modelled population is plotted in Fig. 1. This population shows clear evidence of a mass gap between neutron stars, whose masses go up to ~ 2 solar masses, and black holes, whose masses start at ~ 5 solar masses. This mass gap is a feature of the rapid supernova engine, and reproduces the observed mass gap in neutron star and black hole masses (Özel et al. 2010; Farr et al. 2011) (but see Kreidberg et al. 2012).




Figure 2. *K*-means clustering on the true masses of 400 simulated compact-object binaries. The three clusters perfectly match the actual source subpopulations, with the cluster means shown by magenta Xs.

In the absence of measurement errors, clustering on a subset of observations should be straightforward, and we demonstrate the feasibility of such clustering in Fig. 2. Here, we have chosen 400 merging compact binaries from the population of Fig. 1. The binaries were randomly drawn with a draw probability of NS–NS, NS–BH, and BH–BH binaries set to 25 per cent, 25 per cent, and 50 per cent, respectively¹; the actual population of 400 selected systems has 23 per cent, 26 per cent, and 52 per cent of binaries of the three respective types. These should be interpreted as fractional rates in the observed population, as we do not model selection effects here (cf. Mandel, Farr & Gair 2016). Our goal is to extract these arbitrarily chosen relative contributions of the three subpopulations through clustering.

We performed *K*-means clustering on the exact mass parameters of the 400 binaries. For clustering, we used the mass ratio $q \equiv m_2/m_1$ ² and the chirp mass $M_c \equiv (m_1 + m_2)\eta^{3/5}$, where η is the symmetric mass ratio $\eta \equiv q/(1 + q)^2$. The chirp mass is chosen because it determines the gravitational-wave frequency evolution at the lowest order and is therefore the best measured mass combination (see below). For clustering purposes, we use a logarithmic coordinate on the chirp mass, $\ln(M_c/M_\odot - 0.8)$. Simple *K*-means clustering (MacQueen 1967), which assigns each observation to a cluster with the closest mean, proves adequate for perfect classification on the true source parameters: Every binary in Fig. 2 is correctly assigned to the right cluster.

3 MEASUREMENT UNCERTAINTY

In practice, inference on gravitational-wave signals permits only a limited accuracy of parameter estimation. These limitations are due to significant correlations in the occasionally multimodal parameter space of 15 or more parameters, including component masses and spins, as well as the binary’s sky location and orientation.

¹ These fractions represented an ad hoc choice, not based on the population synthesis model.

² From here on, m_2 is the smaller companion mass; in Fig. 1, it is the mass of the remnant of the secondary star – the star which initially had a lower mass, but could end up as a more massive compact remnant at the end of binary evolution.

Approximate techniques pioneered more than 20 years ago have demonstrated that the chirp mass is a relatively well measured parameter for systems with a total mass of a few tens of solar masses or less, but other mass combinations, such as the mass ratio, can only be relatively poorly constrained (Cutler & Flanagan 1994; Poisson & Will 1995). More recently, Bayesian techniques have been used to directly measure the posterior probability density functions (PDFs) of the signal parameters, given the observed noisy data (Aasi et al. 2013; Abbott et al. 2016c). These techniques, encoded in the `LALINFERENCE` parameter estimation pipeline (Veitch et al. 2015a), have been used to constrain the accuracy of parameter estimation on NS–NS, NS–BH, and BH–BH binaries in a variety of realistic contexts (e.g. Vitale et al. 2014; Littenberg et al. 2015; Mandel et al. 2015; Veitch, Pürrer & Mandel 2015b; Farr et al. 2016; Haster et al. 2016).

Here, we use these earlier results to generate mock posterior PDFs marginalized over all parameters other than m_1 and m_2 . We generate posterior samples in $\{\text{chirp mass } M_c, \text{symmetric mass ratio } \eta\}$ parameter space, given true values (M_c^T, η^T) , as follows:

$$\begin{aligned} M_c &= M_c^T \left[1 + \alpha \frac{12}{\rho} (r_0 + \mathbf{r}) \right], \\ \eta &= \eta^T \left[1 + 0.03 \frac{12}{\rho} (r'_0 + \mathbf{r}') \right]. \end{aligned} \quad (1)$$

Here, r_0 and r'_0 are random numbers drawn from the standard normal distribution, and the corresponding terms encapsulate the shift in the mean of the posterior relative to the true value, while \mathbf{r} and \mathbf{r}' are independent and identically distributed arrays of such random numbers and represent the spread of the posterior. The measurement uncertainty scales inversely with the signal-to-noise ratio ρ , which is drawn from the distribution $p(\rho) \propto \rho^{-4}$, which holds for isotropically distributed sources in a static universe, subject to the threshold $\rho \geq 8$ for detection. The scaling α is motivated by analyses of mock data with the `LALINFERENCE` pipeline (e.g. Littenberg et al. 2015; Mandel et al. 2015) and includes the impact of correlation with parameters describing arbitrary remnant spins; $\alpha = 0.01, 0.03$, and 0.1 when $\eta^T \geq 0.1$, $0.1 > \eta^T \geq 0.05$, and $0.05 > \eta^T$, respectively. Only posterior samples with $0.25 \geq \eta \geq 0.01$ are kept; no a priori cuts on individual masses are assumed, making this an intentionally somewhat conservative estimate of measurement uncertainty.

For each of the 400 mock binaries in our catalogue shown in Fig. 2, we generate between 500 and 2000 posterior samples in $\{m_1, m_2\}$ space, consistent with the typical posterior PDFs produced by `LALINFERENCE` (e.g. Aasi et al. 2013). In Fig. 3, we overplot 500 posterior samples from each of the 400 measured events. Each posterior distribution exhibits a typical ‘banana’-like shape, following contours of roughly constant chirp mass but spanning a range of values of the symmetric mass ratio. The actual size of the posterior depends on the parameter values of the event and its simulated signal-to-noise ratio, while the distribution is randomly shifted relative to the true value so that the true value has a uniform probability of falling into every quantile of the posterior. The combined posterior distributions appear to show an absence of a gap in between NS–NS and NS–BH binaries, and some overlap between NS–BH and BH–BH binaries; meanwhile, due to the lower merger density of higher mass BH–BH binaries in our model, gaps appear at higher masses in the $\{m_1, m_2\}$ distribution.

As an indication of the difficulty of clustering on the observed population suffering from measurement uncertainties, we can apply the *K*-means clustering procedure described in the previous section to the full bag of 400×500 posterior samples plotted in Fig. 3.




Figure 3. 500 posterior samples from each of 400 binaries in the catalogue are placed on the same plot to demonstrate the impact of measurement uncertainty; samples are coloured based on the binary type of the source they are associated with, as in Fig. 1.




Figure 4. *K*-means clustering on a ‘bag’ of posterior samples (500 samples from each of 400 binaries). Some misclassification is evident, and the cluster means, denoted with magenta Xs, no longer correspond to the true subpopulation clusters (cf. Fig. 2).

The result of this attempt is shown in Fig. 4. The very large extent of the posteriors in the q direction makes it difficult to estimate the true locations of the clusters, as evidenced by the shifting of the cluster means in mass ratio relative to their values in Fig. 2. It is still possible to cluster on the chirp mass, however, since it is relatively accurately measured.

However, despite the use of the logarithmic chirp mass coordinate $\ln(M_c/M_\odot - 0.8)$ to aid clustering, some of the BH–BH binaries are misidentified by being associated with the NS–BH cluster. This is most easily seen in Fig. 5, in which we plot each observation at its true mass parameters, but colour it in based on the average cluster association of all corresponding posterior samples. This figure shows yellow (classified as NS–BH) samples in the top-right ‘BH–BH’ cluster (see also the rightmost yellow strip of Fig. 4). Consequently, the fractions of systems in the clusters approximately associated with NS–NS, NS–BH, and BH–BH populations are 23 per cent, 30 per cent, and 47 per cent, respectively. The misclassification errors now exceed those expected from Poisson (trinomial) statistics




Figure 5. K -means clustering on the same bag of posterior samples as in Fig. 4. All 400 observations are displayed at their true parameter values, but the colour reflects the mean cluster association of all posterior samples corresponding to each observation.

for 400 objects, and the classification becomes increasingly poor as the number of observations is reduced.

Of course, the approach described above is flawed because it fails to take advantage of all available information. We lose key information by putting all posterior samples into a single ‘bag’ and ignoring which observation each sample corresponds to. For example, this means that we do not make use of the insight that some posteriors are very broad (and therefore not very useful for clustering), while others correspond to very precise measurements. The right approach must build a hierarchical model out of the full observed population (e.g. Hogg, Myers & Bovy 2010; Mandel 2010; Bovy, Hogg & Roweis 2011; Farr et al. 2015), accounting for the individual measurement uncertainties, and search for subpopulations in this reconstructed population. One possible implementation is described in the following section.

4 DISTRIBUTION INFERENCE AND CLUSTERING

We separate the problem into two parts: hierarchical modelling of the mass distribution based on a finite number of limited-accuracy observations and clustering based on the inferred mass distribution. There are many possible ways to parametrize the mass distribution model. We follow Foreman-Mackey, Hogg & Morton (2014) and Abbott et al. (in preparation) in choosing a piecewise-constant two-dimensional distribution, i.e. we divide the mass space into rectangular bins and model the fraction of systems n_k within each bin $k \in [1, K_{\text{bins}}]$. In this case, we bin in $\ln m_1 \times \ln m_2$ space, with square bins in log space. We cover the range of component masses from $m = 1$ to $181 M_\odot$ with a total of $K_{\text{bins}} = 15 \times 15 = 225$ bins.

When N independent observations are available, each represented by a data set $d^{(i)}$, the posterior PDF on the distribution across the bins $\mathbf{n} \equiv \{n_k\}$ is given by (Mandel 2010)

$$p(\mathbf{n}) \propto \pi(\mathbf{n}) \prod_{i=1}^N p(d^{(i)}|\mathbf{n}). \quad (2)$$

Here

$$p(d^{(i)}|\mathbf{n}) = \int p(d^{(i)}|m_1^{(i)}, m_2^{(i)}) p(m_1^{(i)}, m_2^{(i)}|\mathbf{n}) dm_1^{(i)} dm_2^{(i)}, \quad (3)$$

where $p(d^{(i)}|m_1^{(i)}, m_2^{(i)})$ is the likelihood of observing the data $d^{(i)}$, given the specified masses (Veitch et al. 2015a) and $p(m_1^{(i)}, m_2^{(i)}|\mathbf{n}) = n(m_1^{(i)}, m_2^{(i)})$ is the number density n_k for the appropriate bin k into which these masses fall. In practice, we can replace the preceding integral over the likelihood function by a sum over the available posterior samples, appropriately reweighted by the prior used for individual event analyses (see Mandel 2010; Mandel et al. 2016, for details). Finally, $\pi(\mathbf{n})$ is the prior probability distribution on the fractions within the K_{bins} bins. Our prior is a stationary Gaussian process with a squared-exponential kernel, described in detail in Abbott et al. (in preparation). This prior provides a crucial regularization, favouring a smooth distribution when the data are sparse,³ but allows the posterior to converge to the expected frequentist multinomial distribution when N is large and the measurements are precise.

We compute the inferred distribution according to this hierarchical model from our mock data. Fig. 6 shows the posterior mean of the population density in each of the mass bins inferred from the full set of 400 observations, as well as from smaller randomly drawn subsets, to illustrate the gradual evolution of the accuracy of the inferred posterior. Distinct NS–NS, NS–BH, and BH–BH clusters clearly appear around 40–80 observations, consistently with the estimated requirement of ~ 60 observations made by Mandel et al. (2015).

In order to identify specific clusters, we use a water-filling algorithm on the mean estimates of the population density in each bin (see e.g. Nielsen & Nock 2008; Van & Pham-Gia 2010; Applegate et al. 2011, for other proposed approaches to distributional clustering). We gradually flood the posterior landscape until only three clusters stand above the water level over the $m_1 \geq m_2$ half of the plane. Clusters here are defined as sets of bins such that all elements of a cluster are connected through shared edges, but such connections do not exist between distinct clusters. Some of the posterior ends up in the underwater bins; the clustering is deemed successful only when underwater bins account for no more than a few per cent of the posterior. This happens starting with $N = 80$ for the plots in Fig. 6. As an example, Fig. 7 shows the results of applying the water-filling clustering strategy to the distribution inferred from $N = 400$ observations (mirrored across $m_1 = m_2$ for plotting). In this case, the NS–NS, NS–BH, and BH–BH subpopulations contain 23 per cent, 25 per cent, and 51 per cent of the population, respectively, while less than 2 per cent of the posterior is underwater.

In general, the appropriate number of clusters does not need to be assumed in advance but should be chosen from the data during the water filling stage. Specifically, the amount of water used for flooding can be optimized against the flooded area. Flooding should continue only while the flooded area grows rapidly with a modest increase in the posterior volume (the amount of water used for flooding), with the remaining above-water areas identified as clusters.

We can obtain estimates of the statistical uncertainty on the inferred posterior fraction in each cluster by taking advantage of the full PDFs on the fractional mass distribution within each bin. We use the cluster boundaries provided by the water-filling clustering algorithm and compute the posterior on the total mass density within each cluster identified with the NS–NS, NS–BH, and BH–BH subpopulations. To be precise, given the posterior over all \mathbf{n} , we

³ If sharp edges are expected in the distribution, it would be preferable to use an alternative prior choice that does not disfavour such features.




Figure 6. Mean density inferred across mass space from mock observations using a binned distribution model with a Gaussian process prior for $N = 10, 20, 40$ (top row, left-hand to right-hand side) and 80, 160, 400 (bottom row, left-hand to right-hand side) observations.




Figure 7. Water-filling clustering on the mean estimates of the population fraction in each bin, as inferred from 400 mock observations.

simply add the individual posteriors on the sums of those n_k which fall into a particular cluster; we do not account for the uncertainty in the cluster boundaries when computing these cluster fraction posteriors. The triangle plots for the cluster fraction posteriors are shown in Fig. 8.⁴ For both 80 (left-hand panels) and 400 (right-hand

⁴ For this figure, we associated each underwater bin with a neighbouring cluster; this does not impact the results other than ensuring that the three fractions sum to 1.

panels) observations, the uncertainty in the inferred fraction of each subpopulation is within the expected fluctuation in random-draw statistics from a trinomial distribution, as predicted by Mandel et al. (2015).

We have presented a practical technique for clustering observations suffering from a significant measurement uncertainty. We demonstrated its functionality on the mass parameter space and showed that a realistic population of merging compact binaries could be accurately clustered into NS–NS, NS–BH, and BH–BH subpopulations. The number of observations required for accurate clustering will depend on how well separated the true subpopulations are, on the actual fractions of events in each subpopulation, and on the size of measurement uncertainties. Our example indicates that ~ 20 observations per subpopulation are more than sufficient for accurate clustering on the modelled population. We have confirmed that this number of observations per subpopulation is sufficient for accurate clustering even when the ratio between the numbers of events in different subpopulations is more extreme, e.g. 1: 5: 50 rather than 1: 1: 2. With sufficient observations, it should be possible to use this technique to cluster on any population with multiple modes separated by lower density regions in parameter space (gaps), even if the measurement uncertainty on individual observations is larger than the width of the gaps. It is straightforward, though computationally expensive, to extend this technique to higher dimensional analyses, e.g. to include spin information along with mass information, which could help to distinguish isolated and dynamical formation channels for binary black holes (Abbott et al. 2016a,e).

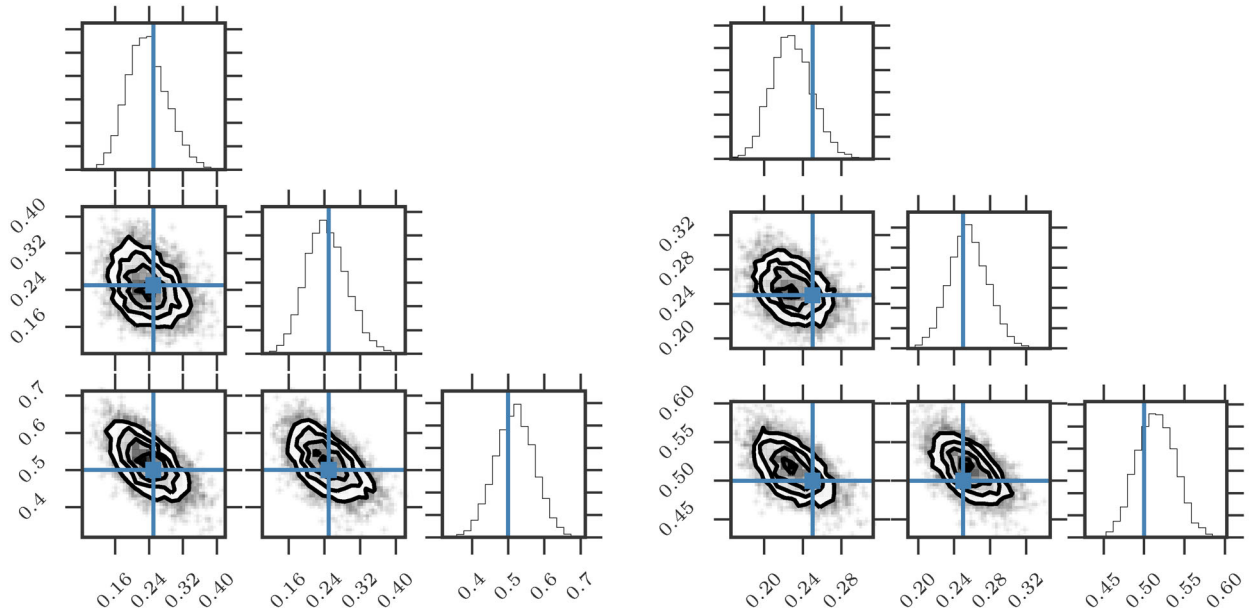


Figure 8. Triangle plots for the posteriors on the inferred fraction of events in each of the NS–NS, NS–BH, and BH–BH subpopulations (ordering from the left-hand to right-hand side, and top to bottom); blue lines denote the fractions used to randomly draw the events being clustered. Left-hand panels: the posterior after 80 mock observations. Right-hand panels: the posterior after 400 mock observations.

ACKNOWLEDGEMENTS

IM and WMF acknowledge partial financial support from STFC.

REFERENCES

- Aasi J. et al., 2013, *Phys. Rev. D*, 88, 062001
Aasi J. et al., 2015, *Class. Quantum Gravity*, 32, 074001
Abadie J. et al., 2010, *Class. Quantum Gravity*, 27, 173001
Abbott B. P. et al., 2016a, *Phys. Rev. X*, 6, 041015
Abbott B. P. et al., 2016b, *Phys. Rev. Lett.*, 116, 061102
Abbott B. P. et al., 2016c, *Phys. Rev. Lett.*, 116, 241102
Abbott B. P. et al., 2016d, *Phys. Rev. Lett.*, 116, 241103
Abbott B. P. et al., 2016e, *ApJ*, 818, L22
Applegate D., Dasu T., Krishnan S., Urbanek S., 2011, in *Proc. 17th ACM SIGKDD Int. Conf., Unsupervised Clustering of Multidimensional Distributions Using Earth Mover Distance*. p. 636, doi:10.1145/2020408.2020508
Bartos I., Kocsis B., Haiman Z., Márka S., 2016, preprint ([arXiv:1602.03831](https://arxiv.org/abs/1602.03831))
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., Wiktorowicz G., Fryer C. L., Holz D. E., Kalogera V., 2012, *ApJ*, 757, 91
Belczynski K., Holz D. E., Bulik T., O’Shaughnessy R., 2016, *Nature*, 534, 512
Bird S., Cholis I., Muñoz J. B., Ali-Haïmoud Y., Kamionkowski M., Kovetz E. D., Racanelli A., Riess A. G., 2016, *Phys. Rev. Lett.*, 116, 201301
Bovy J., Hogg D. W., Roweis S. T., 2011, *Ann. Appl. Stat.*, 5, 1657
Bulik T., Belczyński K., 2003, *ApJ*, 589, L37
Cutler C., Flanagan É. E., 1994, *Phys. Rev. D*, 49, 2658
de Mink S. E., Mandel I., 2016, *MNRAS*, 460, 3545
de Mink S. E., Pols O. R., Hilditch R. W., 2007, *A&A*, 467, 1181
Dominik M., Belczynski K., Fryer C., Holz D. E., Berti E., Bulik T., Mandel I., O’Shaughnessy R., 2012, *ApJ*, 759, 52
Dominik M. et al., 2015, *ApJ*, 806, 263
Eldridge J. J., Stanway E. R., 2016, *MNRAS*, 462, 3302
Farr W. M., Sravan N., Cantrell A., Kreidberg L., Bailyn C. D., Mandel I., Kalogera V., 2011, *ApJ*, 741, 103
Farr W. M., Gair J. R., Mandel I., Cutler C., 2015, *Phys. Rev. D*, 91, 023005
Farr B. et al., 2016, *ApJ*, 825, 116
Foreman-Mackey D., Hogg D. W., Morton T. D., 2014, *ApJ*, 795, 64
Fryer C. L., Belczynski K., Wiktorowicz G., Dominik M., Kalogera V., Holz D. E., 2012, *ApJ*, 749, 91
Haster C.-J., Wang Z., Berry C. P. L., Stevenson S., Veitch J., Mandel I., 2016, *MNRAS*, 457, 4499
Hogg D. W., Myers A. D., Bovy J., 2010, *ApJ*, 725, 2166
Inayoshi K., Kashiyama K., Visbal E., Haiman Z., 2016, *MNRAS*, 461, 2722
Ivanova N. et al., 2013, *A&AR*, 21, 59
Kreidberg L., Bailyn C. D., Farr W. M., Kalogera V., 2012, *ApJ*, 757, 36
Lipunov V. M., Kornilov V., Gorboskoy E., Tiurina N., Balanutsa P., Kuznetsov A., 2016, *New Astron.*, 51, 122
Littenberg T. B., Farr B., Coughlin S., Kalogera V., Holz D. E., 2015, *ApJ*, 807, L24
MacQueen J. B., 1967, in *Le Cam L. M., Neyman J., eds, Berkeley Symp. Math. Statist. Prob., Proceedings of 5th Berkeley Symposium on Mathematical Statistics and Probability*. Univ. California Press, p. 281
Mandel I., 2010, *Phys. Rev. D*, 81, 084029
Mandel I., 2016, *MNRAS*, 456, 578
Mandel I., de Mink S. E., 2016, *MNRAS*, 458, 2634
Mandel I., O’Shaughnessy R., 2010, *Class. Quantum Gravity*, 27, 114007
Mandel I., Haster C.-J., Dominik M., Belczynski K., 2015, *MNRAS*, 450, L85
Mandel I., Farr W. M., Gair J. R., 2016, *Extracting Distribution Parameters from Multiple Uncertain Observations with Selection Biases*. Available at: <https://dcc.ligo.org/LIGO-P1600187/public>
Mapelli M., 2016, *MNRAS*, 459, 3432
Marchant P., Langer N., Podsiadlowski P., Tauris T. M., Moriya T. J., 2016, *A&A*, 588, A50
Mennekens N., Vanbeveren D., 2014, *A&A*, 564, A134
Nielsen F., Nock R., 2008, *Emerging Trends in Visual Computing*. Springer-Verlag, Berlin
O’Shaughnessy R., 2013, *Phys. Rev. D*, 88, 084061
Özel F., Psaltis D., Narayan R., McClintock J. E., 2010, *ApJ*, 725, 1918
Poisson E., Will C. M., 1995, *Phys. Rev. D*, 52, 848

- Postnov K. A., Yungelson L. R., 2014, *Living Rev. Relativ.*, 17, 3
- Repetto S., Nelemans G., 2015, *MNRAS*, 453, 3341
- Rodriguez C. L., Haster C.-J., Chatterjee S., Kalogera V., Rasio F. A., 2016, *ApJ*, 824, L8
- Stevenson S., Ohme F., Fairhurst S., 2015, *ApJ*, 810, 58
- Stone N. C., Metzger B. D., Haiman Z., 2016, *MNRAS*, 464, 946
- Van T. V., Pham-Gia T., 2010, *J. Appl. Stat.*, 37, 1891
- Veitch J. et al., 2015a, *Phys. Rev. D*, 91, 042003
- Veitch J., Pürrer M., Mandel I., 2015b, *Phys. Rev. Lett.*, 115, 141101
- Vink J. S., de Koter A., Lamers H. J. G. L. M., 2001, *A&A*, 369, 574
- Vitale S., Lynch R., Veitch J., Raymond V., Sturani R., 2014, *Phys. Rev. Lett.*, 112, 251101

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