



Capano, C., Dent, T., Hanna, C., Hendry, M., Messenger, C., Hu, Y.-M. and [Veitch, J.](#) (2017) Systematic errors in estimation of gravitational-wave candidate significance. *Physical Review D*, 96(8), (doi:[10.1103/PhysRevD.96.082002](https://doi.org/10.1103/PhysRevD.96.082002))

This is the author's final accepted version.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/152680/>

Deposited on: 05 January 2018

Enlighten – Research publications by members of the University of Glasgow
<http://eprints.gla.ac.uk>

Systematic errors in estimation of gravitational-wave candidate significance

C. Capano and T. Dent

Max Planck Institute for Gravitational Physics, Callinstraße 38, D-30167, Hannover, Germany

C. Hanna

The Pennsylvania State University, University Park, PA 16802, USA

M. Hendry and C. Messenger

School of Physics and Astronomy, Kelvin Building, University of Glasgow, Glasgow, G12 8QQ, UK

Y.-M. Hu*

TianQin Research Center for Gravitational Physics, Sun Yat-sen University, Zhuhai, 519082, China
School of Physics and Astronomy, Kelvin Building, University of Glasgow, Glasgow, G12 8QQ, UK and
Max Planck Institute for Gravitational Physics, Callinstraße 38, D-30167, Hannover, Germany

J. Veitch

School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham, B15 2TT, UK and
School of Physics and Astronomy, Kelvin Building, University of Glasgow, Glasgow, G12 8QQ, UK

(Dated: August 23, 2017)

(LIGO-P1500247-v1)

The statistical significance of a candidate gravitational-wave (GW) event is crucial to the prospects for a confirmed detection, or for its selection as a candidate for follow-up electromagnetic observation. To determine the significance of a GW candidate, a ranking statistic is evaluated and compared to an empirically-estimated background distribution, yielding a false alarm probability or p -value. The reliability of this background estimate is limited by the number of background samples and by the fact that GW detectors cannot be shielded from signals, making it impossible to identify a pure background data set. Different strategies have been proposed: in one method, *all samples*, including potential signals, are included in the background estimation, whereas in another method, *coincidence removal* is performed in order to exclude possible signals from the estimated background. Here we report on a mock data challenge, performed prior to the first detections of GW signals by Advanced LIGO, to compare these two methods. The all-samples method is found to be self-consistent in terms of the rate of false positive detection claims, but its p -value estimates are systematically conservative and subject to higher variance. Conversely, the coincidence-removal method yields a mean-unbiased estimate of the p -value but sacrifices self-consistency. We provide a simple formula for the uncertainty in estimate significance and compare it to mock data results. Finally, we discuss the use of different methods in claiming the detection of GW signals.

I. INTRODUCTION

The first direct detections of gravitational wave (GW) signals from the mergers of massive black hole binary systems achieved substantial scientific and societal impact when reported in 2016 [1–17]. A key aspect of the data analysis required for these detections, and an essential step in verifying their identity as astrophysical GW signals, was robustly establishing the statistical significance of candidate events in the search for GW from binary mergers.

Extraordinary claims require extraordinary evidence, and the discovery of a previously unseen effect or signal has generally required the significance to exceed a certain threshold. Consequently, in order to reject confidently the null hypothesis that one’s dataset contains *only* noise and claim the presence of a signal, the probability of obtaining this dataset when the null hypothesis is true, i.e. the p -value, must be no larger than a pre-determined threshold. The threshold for discovery

is often set at “5 sigma”, i.e. the probability of obtaining a value 5 or more standard deviations from mean of a Gaussian distribution [18], as for instance in the detection of the Higgs boson [19] and B-modes in the polarization of the cosmic microwave background (CMB) [20]. Less stringent significance thresholds have typically been applied when presenting *evidence* for a possible new effect, and for applications where the rate of false positives is not required to be extremely low – for example in selecting candidate GW events for follow-up observations over the electromagnetic spectrum [21].

Ground-based detector GW data analysis presents several unique challenges that complicate the calculation and interpretation of the significance. Due to the noisy local environments at each detector, the simultaneous observation of candidate events (or ‘triggers’) by multiple detectors, with consistent signal characteristics observed at each detector, is a necessary condition. Moreover, the low signal-to-noise ratio (SNR) expected means that matched-filtering of the data with signal templates is required, but similarity among the templates used might lead to correlation of the triggers’ SNR. Also, no *a priori* information about the background distribution is available, so this has to be estimated from the observations themselves,

* huyiming@mail.sysu.edu.cn

and the possible contamination by a real GW signal cannot simply be spotted and removed. The expected rate of GW signals, prior to any positive detection, is also uncertain by several orders of magnitude [22].

To address these issues, analysis methods have been developed to discover potential GW candidates and assign significance to them [23–26]. These methods use similar ideas, albeit with quite different specific implementations. More importantly, however, the methods encapsulate two different philosophical viewpoints about the correct way to calculate the significance of a candidate event. In this manuscript we describe a mock data challenge (MDC), motivated by seeking to resolve the conflict between these two viewpoints, which emerged in the “Blind Injection” exercise undertaken by LIGO and Virgo in 2010 [27]. We demonstrate quantitatively the pros and cons of both viewpoints for quantifying the significance of GW candidates via a *false alarm probability* or *p*-value.

This MDC was designed, executed and concluded before the first GW detections and its results informed the confident detections of GW150914, GW151226 and GW170104, as well as the evaluation of the less significant candidate event LVT151012. With detections of binary neutron star (BNS) and neutron star black hole (NSBH) events also expected in the next few years, our MDC will provide the GW community with clear guidelines for computing and interpreting the false alarm probabilities of future GW detection candidates as well as justifying the choices made for previous results. A detailed account of the setup, results and conclusions from the MDC can be found in [28]; here we summarize the essential features of the study and describe the main results relevant to GW detection.

II. THE MOCK DATA CHALLENGE

A. Rationale

In order to assess the significance of a GW trigger, one must understand the statistical properties of the distribution of background noise. Unfortunately one cannot simply shield the detectors from GW signals and thus measure a background that is guaranteed to be ‘signal-free’. However, various methods have been developed to approximate such a background [23–25].

Ground-based GW interferometers constantly generate triggers, but only those triggers that excite a high response in more than one detector simultaneously, in a way consistent with an astrophysical signal, are regarded as viable GW candidates. These are then known as coincident or *zero-lag* triggers. The controversy around background estimation focuses on how to deal with these zero-lag triggers.

One approach is simply to include zero-lag triggers in the background estimation – consistent with the argument that no trigger will be classified as a confirmed signal until it passes a pre-determined threshold, so no trigger should be removed prior to that confirmation. This approach might result in actual GW signals contaminating the estimation of the noise back-

ground, but in principle this should be a *conservative* contamination – i.e. any trigger identified as passing the threshold would also have done so with an uncontaminated background.

The other approach is to use only single-interferometer (IFO) triggers to estimate the noise background, having first removed the zero-lag triggers – arguing that the latter are intrinsically different as they are (the only) viable GW candidates and thus it is inappropriate to consider them when estimating the background distribution. In principle, removing background-induced zero-lag triggers should result in an unbiased estimate of the significance of a GW candidate.

B. MDC Setup and Definitions

To better understand, and hopefully resolve, the potential impact of zero-lag trigger removal on the estimation of significance, we therefore sought to carry out a MDC in which all significance calculations using the mock data would be done using both ‘coincidence-removal’ and ‘all-coincidence’ methods. Since our focus was on this one issue, we recognised that an end-to-end simulation of GW detection would involve lots of technical details not directly linked to the specific question of ‘removal’ versus ‘non-removal’. Thus we sought to make the MDC as simple as possible.

In generating our mock data, each trigger was labeled with two quantities: a ranking statistic (or SNR) ρ_i for the i^{th} IFO, and an arrival time. A background trigger will occur at each IFO independently while an astronomical signal would trigger responses dependent on the geometry of the detectors, as well as that of the source. The frequency of background triggers and astronomical signals are each controlled by a rate parameter. To mimic choices typically used in LIGO-Virgo analysis [3], we set a threshold of $\rho_i^{\text{th}} = 5.5$ for a trigger to be registered in a single IFO.

A time window based on the light travel time across the Earth was used to identify all coincident triggers among different detectors. The SNR, ρ , of the coincident trigger is the root sum square of the SNR, ρ_i , for the individual detectors. This ranking statistic ρ for each coincident pair of triggers is used to calculate the significance of the zero-lag events.

The parameters of the MDC were chosen so that each *realisation* contained $\sim 10^4$ background triggers in each IFO, and statistically ~ 10 zero-lag triggers were expected to originate from random coincidences in the background distribution. Each *experiment* then contained 10^5 realisations. This choice allowed us to study different interesting regimes, while keeping the computational burden of the MDC practically feasible.

In total we designed 14 experiments with background distributions of different levels of complexity, which we label respectively ‘simple’, ‘realistic’ and ‘extreme’; the astronomical signal rate in these experiments ranges from zero (i.e. no astronomical events at all) through low (an expected 0.001 – 0.1 coincidences per realisation), medium (~ 0.5 coincidences per realisation) to high (~ 3 coincidences per realisation). Crucially, the fact that the background distribution was known to the MDC designers made it possible to calculate the exact de-

pendence of the false alarm probability (FAP) on the ρ value for each experiment.

We injected BNS signals with anticipated signal-to-noise ratio (SNR) calculated from the inspiral part of the signal only, assuming random sky locations and distances, uniform in co-moving volume, up to a cutoff. The actual measured SNR was randomly assigned based on the simulated signal. The distributions for both signals and noise were taken to be rapidly decreasing at high SNR, but the distribution of astronomical signals had a shallower (power-law) slope: thus signals ‘stand out’ in the region with larger ρ .

Of course some astronomical events might cause a trigger in one detector only, as the SNR in one or more other detectors may not be high enough to exceed the threshold. Since in this MDC we limit our investigation only to the impact of removal or non-removal on the *loudest* (i.e. highest SNR ρ) event, which is not expected to originate from such a signal-noise hybrid, we can safely ignore any adverse impact of this scenario.

C. MDC participants

We define the significance of the loudest coincidence in terms of its FAP. That is, the probability of observing a coincidence with ρ higher than the recorded loudest zero-lag event.

Three different algorithms for estimating the background distribution from the observed triggers were investigated in the MDC.

1. The ‘time slide’ algorithm adapted from the all-sky LIGO-Virgo search pipeline [23, 24] shifts one detector’s triggers by a certain time interval so that, in the shifted data, no new coincidence could be associated with a real astronomical event. Thus the statistical properties of the background may be estimated. A false alarm rate (FAR) is calculated for a coincidence with a particular value of ρ by evaluating the frequency of coincidences per analysis time that have as high or higher a value of ρ . The inverse FAR (IFAR) algorithm is used to estimate the probability of obtaining one or more coincidences with higher SNR. By exhausting all allowable time shifts, one can estimate both the expected number of background coincidences and the probability for each to exceed a given ρ . The FAP is then calculated by taking account of the number of trials from the expected coincidence number.
2. The ‘all possible coincidences (APC)’ algorithm used a similar strategy to the time slide algorithm, by combining all physically unrelated triggers from different detectors to form an ensemble of all possible coincidences. The FAP was then calculated simply by counting the fraction of coincidences louder than the loudest observed event, with trial factors included. In addition to the FAP, the APC algorithm can also provide an estimate of its associated uncertainty. The smaller the FAP is, the higher the relative uncertainty will be. In this

MDC, the APC algorithm dynamically adjusted the calculation precision, so that the estimation of larger FAP had a higher uncertainty.

3. A modification of the *gstlal* algorithm [25] was used to estimate the FAP. From the observed triggers, their distribution can be estimated for each single IFO individually, and then used to estimate the distribution of ρ . The FAP was then calculated from this extrapolated distribution.

All three participating algorithms were applied with both ‘coincidence removal’ and ‘all samples’ methods, giving six different outputs for each realisation. In this MDC, the three different algorithms set different lower limits on the FAP estimation: the IFAR algorithm adopted a lower limit of 10^{-7} , the *gstlal* algorithm adopted 10^{-6} , and the APC algorithm set no hard lower limit.

D. Estimating the error on the false alarm probability

Essentially, all three algorithms assumed independence of the background distributions across different detectors, and in all cases estimation of the FAP involved counting the fraction of all coincidences with larger SNR than a given ρ . The uncertainty on the calculated FAP thus followed a simple Poisson counting error, scaling with the square root of the number of coincidences with larger SNR. In the limit of small FAP, which is the most interesting region, the relative uncertainty of FAP can be roughly approximated to be inversely proportional to the square root of the number of triggers within one IFO that is louder than ρ_i . A detailed derivation and discussion of this result is available in [28]; see also [29] for related work on uncertainties in GW search background estimation.

III. RESULTS

In the detailed technical paper [28], more comprehensive analysis of the MDC results is carried out; here we focus on general conclusions. We examined the estimated FAPs for all 14 experiments, applying all three algorithms with both methods of treating candidate triggers. We then investigated our results for issues of self-consistency, precision, bias, and detection efficiency. Signal-free experiments were exclusively adopted in subsection III A while conversely only experiments containing astronomical signals were considered in subsection III D. All experiments were included in subsection III B and III C.

A. p - p plots

There were four experiments that contained no astronomical signals. For these *noise only* experiments, a self-consistency test can be performed by drawing a p - p plot. For each nominal value of FAP, p , one can count the fraction within all 10^5 realisations that has a smaller estimated FAP.

For a self-consistent estimator on a pure noise dataset, the fraction of events having lower FAP than p should be equal to p , to within the counting uncertainty.

In Fig. 1a we show the p - p plot for one specific experiment where a clear discrepancy between the results for ‘all samples’ and ‘coincidence removal’ methods is apparent for estimated FAP values of around 10^{-3} or less. The ‘all samples’ results show good self-consistency, while the ‘coincidence removal’ results show a clear tendency to underestimate the FAP, for small FAP values. This feature is universally observed, but will be most severe when the loudest single trigger forms a zero-lag coincidence and is removed – in which case it can be understood in terms of the following heuristic explanation. For the distributions we consider, the loudest trigger in one or other IFO may often be a moderate outlier above the bulk of the noise distribution. If seen in combination with this loudest trigger, almost every single trigger from the other IFO would produce a loud coincidence. Thus, if the loudest trigger is removed, essentially the FAP will be underestimated by $O(1/N)$, where N is the number of triggers in one IFO. For n trials (where $n \ll N$) one can expect this scenario to affect $\sim n/N$ of the realisations. Taking a typical value of $n = 10$ and $N = 10^4$, would mean that a fraction of the loudest 10^{-3} realisations would be affected by this underestimation – consistent with our results. Although the probability that the loudest trigger in one IFO forms a random coincidence with a noise trigger in the other is small, it is *larger* than the factor by which the FAP could be underestimated as a consequence of its removal.

We also investigated if this observed breakdown of self-consistency could be associated with the timing of the triggers. We designed a simple replica of the MDC in which only SNR was involved, so that the ‘pairing’ of coincident triggers was completely random. Both ‘all samples’ and ‘coincidence removal’ methods were used to estimate the FAP in this case. Fig. 1b shows an example where $N = 100$ and $n = 1$. As predicted by our heuristic explanation, the discrepancy of the ‘coincidence removal’ method becomes apparent when $n/N = 10^{-2}$. We conclude, therefore, that the ‘coincidence removal’ method’s loss of self-consistency has its origin in the methodology itself, i.e. the exclusion of candidate samples which substantially affects the background estimate in a small number of cases.

B. Direct comparison

We also compared estimated FAP values with exact FAP values using all experiments. All three algorithms show similar behaviour; in each case, however, different results were obtained for the two different methods of treating candidate triggers. One typical result is shown in Fig. 2, where the general pattern displayed is that the smaller the FAP, the larger the relative uncertainty in its estimated value. The spread of relative uncertainties is consistent with predictions.

It can be seen that, in the small FAP limit, by removing all zero-lag coincidences the FAP values are much more likely to be underestimated. Under the condition that the exact FAP is small, but bounded by the lower limit of zero,

any underestimation can only trivially decrease the expected value, while any potential overestimation would contribute disproportionately. The conditional probability distribution $p(\text{FAP}_{\text{estimate}} | \text{FAP}_{\text{true}} \ll 1, \text{Removal})$ would thus naturally become skewed, and it is very likely that such an estimator could report an estimate of, say, 10^{-5} , when the exact FAP is actually 10^{-4} .

C. Box plot

In order to investigate the bias of the FAP estimates, using all experiments we selected realisations with exact FAP values in the range $10^{-4} - 10^{-3}$. Fig. 3, shows box plots for the ratio of estimated FAP and exact FAP. We can see that for the ‘all samples’ method, for all algorithms, the median estimators are essentially unbiased within sampling uncertainty while the mean values are generally overestimated. On the other hand, results from the ‘coincidence removal’ method show unbiased mean values of the FAP estimator, but median values tend to be underestimated. This result shows vividly how an unbiased estimator of FAP, in the small FAP limit will nonetheless lose self-consistency for the majority of realisations.

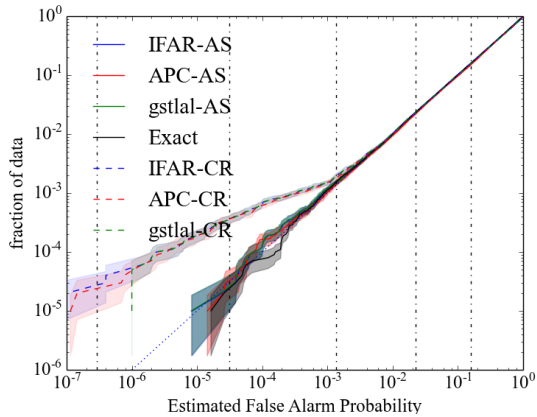
D. Receiver Operating Characteristic

Finally, using only those experiments containing astronomical signals, we plotted changed the ‘‘plotted’’ to ‘‘obtained’’ Receiver Operating Characteristic (ROC) plots to compare the detection efficiency of our different algorithms using both methods of treating candidate triggers. In these plots the false positive rate (FPR) is the fraction of noise realisations with FAP smaller than a certain threshold, and true positive rate (TPR) is this same fraction for astronomical signals. The closer a ROC curve is located to the top left corner of the diagram, the higher is the detection efficiency of the estimator.

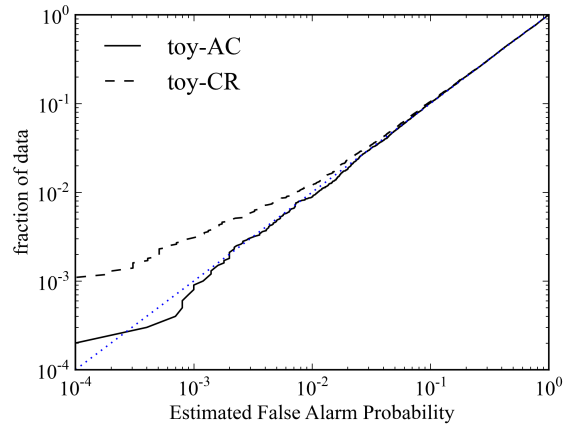
Our results showed that all algorithms gave ROC curves with similar behaviour, and for each algorithm the difference between the ROC curves for the two methods of treating candidate triggers lay mostly within the sampling uncertainty. There were several specific experiments where the ‘all samples’ method was found to have higher TPR for the smallest FPR, and the difference was larger than the numerical uncertainty. Further investigation indicated that these discrepancies between the two methods occurred at similar FAP values as the discrepancies that arose in the p - p plots. By falsely assigning too much background noise with a low FAP, the detection of signals with the ‘coincidence removal’ method is indeed rendered more difficult. However, this feature is not universally observed for all experiments.

IV. SUMMARY AND CONCLUSIONS

We have set up and performed a MDC to investigate the impact of removing zero-lag coincidences on the estimation of FAP for GW candidates. 14 experiments were conducted,

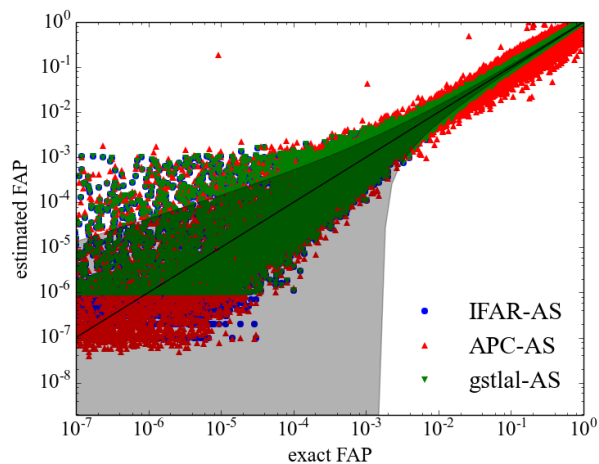


(a) The p - p plot for one experiment. The vertical dashed lines correspond to different numbers of Gaussian standard deviations from 1σ through 5σ (right to left).

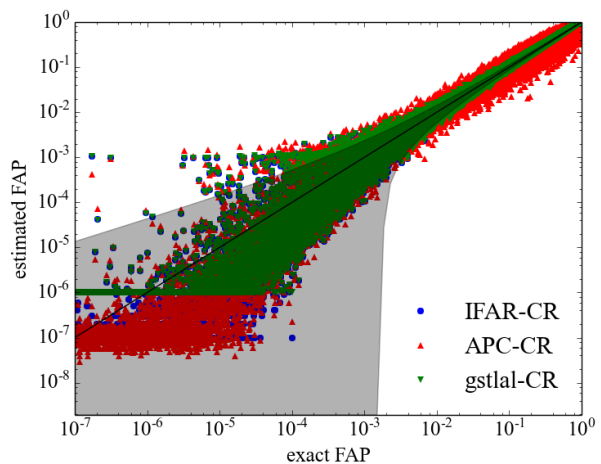


(b) The p - p plot for a toy ‘MDC’ in which SNR and associated ‘coincidence’ are assigned completely randomly.

FIG. 1: The p - p plots for one experiment with no astronomical signals (left) and for a toy model (right). The x -axis shows the estimated FAP value, and the y -axis shows the fraction of events over all realisations with estimated FAP equal to or smaller than that value. A self-consistent estimate should lie around the diagonal line. The solid lines represent the results for the different algorithms with the ‘all samples’ (AS) method, while the dashed lines show results with the ‘coincidence removal’ (CR) method.



(a) ‘All samples’ comparison.



(b) ‘Coincidence removal’ comparison.

FIG. 2: Directly comparing FAP estimates with the exact FAP for one typical experiment. Shaded regions represent predicted uncertainty. Different methods give comparable results with visible difference of lower boundary due to *a priori* choice. APC method shows larger scatter in larger FAP due to lower accuracy used, which does not affect the result of interesting small FAP regions. Left panel shows results for ‘all samples’ (AS) method and right panel for ‘coincidence removal’ (CR).

each comprising 10^5 realisations, covering a wide range of background distribution complexity as well as widely different astronomical signal rates. Not all experiments were designed to represent our understanding of the signals and backgrounds expected for the current generation of ground-based GW detectors; some extreme cases were included to test the robustness of current FAP estimation methods. Three different algorithms for background estimation were used, each using two methods of treating candidate triggers, namely ‘coincidence removal’ and ‘all samples’. Throughout all experiments, the three different algorithms showed substantially

similar behaviour; obvious differences arose only due to different choices of FAP cutoff and calculation accuracy.

However, the two different methods of treating candidate triggers showed clear differences in the estimated FAP values – with the ‘coincidence removal’ method generally having an unbiased mean value while the ‘all samples’ method was self-consistent. We demonstrated that, for the most interesting regime of small FAP values, self-consistency and unbiasedness for the mean *cannot* be achieved simultaneously.

We recommend that, for the first detections of GW from a previously unobserved source, particularly when signal rates

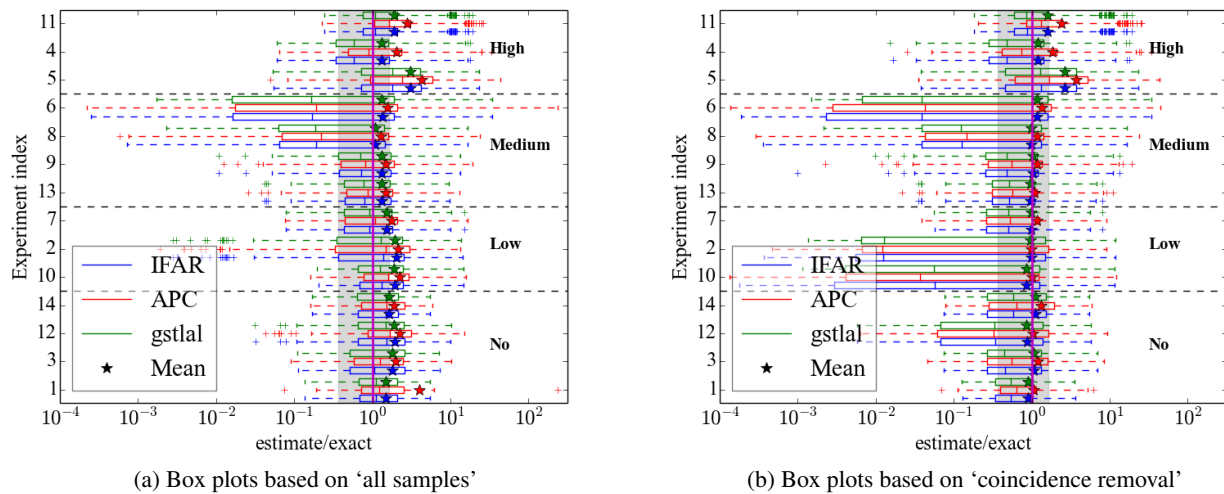


FIG. 3: Box plots of ratios between estimated and exact FAP values, based on realisations with exact FAPs $\in [10^{-4}, 10^{-3}]$. The 14 experiments are ordered by descending astrophysical signal rates. The shaded region indicates the expected uncertainty in the ratio. The middle box indicates the interquartile range, the thin vertical line within the box is the median, and the star represents the mean.

are highly uncertain, a strict threshold should be imposed on the self-consistent FAP calculated via the ‘all samples’ method. Conversely, the mean-unbiased nature of the ‘coincidence removal’ method means its estimate of the FAP as a function of the search ranking statistic is more informative about the background distribution.

With the expected increase in sensitivity of the Advanced detector network [17], the rate of true signals is expected to increase over time. With a higher rate of astronomical events, *a priori* the loudest candidate event is correspondingly more likely to be signal than noise. Therefore, in the longer run, we would expect to place more importance on mean unbiasedness than on self-consistency, which favours the ‘coincidence removal’ method in searches where the rate of detectable signals is known to be high. However, for cases where > 1 loud signals occur per experiment, we anticipate that Bayesian methods based on modelling signal and noise distributions simultaneously, with a well-defined signal prior [26] will be more suitable.

We also computed a heuristic estimate of the relative uncertainty in the FAP value: the square root of the number of triggers within one IFO that are louder than ρ_i . This predicted

uncertainty is roughly consistent with the numerical scatter observed in our realisations. Moreover, the deviation from self-consistency affects the FAP especially strongly for values less than or equal to n/N , where n is the expected number of coincidences and N is the total number of triggers in one IFO. We note that, due to the expedient choice of parameters in our simulations, our quantitative numerical conclusions are strictly only valid in the context of the MDC, although they do offer an instructive indication of the expected behaviour for more realistic cases.

ACKNOWLEDGMENTS

The authors have benefited from discussions with many LIGO Scientific Collaboration (LSC) and Virgo collaboration members and ex-members including Drew Keppel and Matthew West. We acknowledge the use of LIGO Data Grid computing facilities for the generation and analysis of the MDC data. C. M. is supported by a Glasgow University Lord Kelvin Adam Smith Fellowship and the Science and Technology Research Council (STFC) grant No. ST/L000946/1. M. H. is supported by STFC grant No. ST/L000946/1. J. V. is supported by STFC grant No. ST/K005014/1.

-
- [1] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, and et al., *Physical Review Letters* **116**, 061102 (2016), arXiv:1602.03837 [gr-qc].
 - [2] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Phys. Rev. Lett.* **116**, 131103 (2016), arXiv:1602.03838 [gr-qc].
 - [3] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Phys. Rev.* **D93**, 122003 (2016), arXiv:1602.03839 [gr-qc].
 - [4] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Phys. Rev. Lett.* **116**, 241102 (2016), arXiv:1602.03840 [gr-qc].
 - [5] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Phys. Rev. Lett.* **116**, 221101 (2016), arXiv:1602.03841 [gr-qc].
 - [6] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Astrophys. J.* **833**, L1 (2016), arXiv:1602.03842 [astro-ph.HE].
 - [7] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Phys. Rev.* **D93**, 122004 (2016), [Addendum: *Phys. Rev.* **D94**, no.6, 069903(2016)], arXiv:1602.03843 [gr-qc].
 - [8] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Class. Quant. Grav.* **33**, 134001 (2016), arXiv:1602.03844 [gr-qc].
 - [9] B. P. Abbott *et al.* (LIGO Scientific), *Phys. Rev.* **D95**, 062003

- (2017), arXiv:1602.03845 [gr-qc].
- [10] B. P. Abbott *et al.* (Virgo, LIGO Scientific), *Astrophys. J.* **818**, L22 (2016), arXiv:1602.03846 [astro-ph.HE].
- [11] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, and *et al.*, *Physical Review Letters* **116**, 131102 (2016), arXiv:1602.03847 [gr-qc].
- [12] B. P. Abbott *et al.*, *Physical Review Letters* **116**, 241103 (2016).
- [13] B. P. Abbott *et al.* (The LIGO Scientific Collaboration and the Virgo Collaboration), *ArXiv e-prints* (2016), arXiv:1606.04856.
- [14] G. M. Harry and the LIGO Scientific Collaboration, *Classical and Quantum Gravity* **27**, 084006 (2010).
- [15] J. Aasi *et al.* (LIGO Scientific Collaboration), *Class. Quantum Grav.* **32**, 074001 (2015), arXiv:1411.4547 [gr-qc].
- [16] F. Acernese *et al.* (VIRGO), *Class. Quantum Grav.* **32**, 024001 (2015), arXiv:1408.3978 [gr-qc].
- [17] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), *Living Rev. Relat.* **19**, 1 (2016), arXiv:1304.0670 [gr-qc].
- [18] L. Lyons, (2013), arXiv:1310.1284 [physics.data-an].
- [19] G. Aad, T. Abajyan, B. Abbott, J. Abdallah, S. Abdel Khalek, A. A. Abdelalim, O. Abidinov, R. Aben, B. Abi, M. Abolins, *et al.*, *Physics Letters B* **716**, 1 (2012), arXiv:1207.7214.
- [20] P. A. R. Ade *et al.* (BICEP2 Collaboration), *Phys. Rev. Lett.* **112**, 241101 (2014).
- [21] B. P. Abbott *et al.* (InterPlanetary Network, DES, INTEGRAL, La Silla-QUEST Survey, MWA, Fermi-LAT, J-GEM, DEC, GRAWITA, Pi of the Sky, Fermi GBM, MASTER, Swift, iPTF, VISTA, ASKAP, SkyMapper, PESSTO, TOROS, Pan-STARRS, Virgo, Liverpool Telescope, BOOTES, LIGO Scientific, LOFAR, C2PU, MAXI), *Astrophys. J.* **826**, L13 (2016), arXiv:1602.08492 [astro-ph.HE].
- [22] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Class. Quantum Grav.* **27**, 173001 (2010).
- [23] S. Babak, R. Biswas, P. Brady, D. Brown, K. Cannon, *et al.*, *Phys. Rev. D* **87**, 024033 (2013), arXiv:1208.3491 [gr-qc].
- [24] B. Abbott *et al.* (LIGO Scientific Collaboration), *Phys. Rev. D* **80**, 047101 (2009).
- [25] K. Cannon, C. Hanna, and D. Keppel, *Phys. Rev. D* **88**, 024025 (2013).
- [26] W. M. Farr, J. R. Gair, I. Mandel, and C. Cutler, *Phys. Rev. D* **91**, 023005 (2015), arXiv:1302.5341 [astro-ph.IM].
- [27] J. Abadie *et al.* (LIGO Collaboration, Virgo Collaboration), *Phys. Rev. D* **85**, 082002 (2012), arXiv:1111.7314 [gr-qc].
- [28] C. Capano *et al.*, (2015), LIGO document P1500247-v3, <https://dcc.ligo.org/LIGO-P1500247/public>.
- [29] M. Was, M.-A. Bizouard, V. Brisson, F. Cavalier, M. Davier, *et al.*, *Class. Quantum Grav.* **27**, 015005 (2010), arXiv:0906.2120 [gr-qc].