

Overlapping signals in next-generation gravitational wave observatories: A recipe for selecting the best parameter estimation technique

Tomasz Baka^{1,2}, Harsh Narola^{1,2}, Justin Janquart^{3,4}, Anuradha Samajdar^{1,2},
Tim Dietrich^{5,6} and Chris Van Den Broeck^{1,2}

¹*Institute for Gravitational and Subatomic Physics (GRASP), Utrecht University,
3584 CC Utrecht, The Netherlands*

²*Nikhef, 1098 XG Amsterdam, The Netherlands*

³*Center for Cosmology, Particle Physics and Phenomenology—CP3,
Université Catholique de Louvain, Louvain-La-Neuve, B-1348, Belgium*

⁴*Royal Observatory of Belgium, Avenue Circulaire, 3, 1180 Uccle, Belgium*

⁵*Institut für Physik und Astronomie, Universität Potsdam, Haus 28,
Karl-Liebknecht-Strasse 24/25, 14476, Potsdam, Germany*

⁶*Max Planck Institute for Gravitational Physics (Albert Einstein Institute),
Am Mühlenberg 1, Potsdam, Germany*



(Received 22 July 2025; accepted 12 August 2025; published 3 October 2025)

Third-generation gravitational wave detectors such as Einstein Telescope and Cosmic Explorer will have significantly better sensitivities than current detectors, as well as a wider frequency bandwidth. This will increase the number and duration of the observed signals, leading to many signals overlapping in time. If not adequately accounted for, this can lead to biases in parameter estimation. In this work, we combine the joint parameter estimation method with relative binning to handle full parameter inference on overlapping signals from binary black holes, including precession effects and higher-order mode content. As this method is computationally more expensive than traditional single-signal parameter estimation, we test a prior-informed Fisher matrix and a time-frequency overlap method for estimating expected bias to help us decide when joint parameter estimation is necessary over the simpler methods. We improve upon previous Fisher matrix implementations by including the prior information and performing an optimization routine to better locate the maximum likelihood point, but we still find the method unreliable. The time-frequency method is accurate in 86% of close binary black hole mergers. We end by developing our own method of estimating bias due overlaps, where we reweight the single signal parameter estimation posterior to quantify how much the overlapping signals affect it. We show it has 99% accuracy for zero noise injections (98% in Gaussian noise), at the cost of one additional standard sampling run when joint parameter estimation proves to be necessary.

DOI: [10.1103/8cwp-mxcd](https://doi.org/10.1103/8cwp-mxcd)

I. INTRODUCTION

The continuous upgrade of the Advanced LIGO [1] and Advanced Virgo [2] detectors over the last ten years has made the observation of gravitational waves (GWs) routine, with already more than 200 GW candidate detections reported during the O4 observing runs [3]. The in-depth analysis of the events observed during the previous observing runs [4] has already enabled major advances in fundamental physics [5], astrophysics [6], and cosmology [7], and the increase in number of observed GW signals should further improve the precision of the measurements. Even further upgrades to these second-generation detectors [8], as well as the addition of other detectors to the network—such as KAGRA [9,10] and LIGO-India [11]—will further increase the number of observed signals and the precision of the measurements. Additionally, a next generation of ground-based detectors is

planned, with Einstein Telescope (ET) [12,13] in Europe and Cosmic Explorer (CE) [14] in the United States; these are referred to as third-generation (3G) observatories. They are expected to lead to a steep change in GW observations, both in terms of sensitivity and increased frequency bandwidth [15]. As a result, it will become common for signals to overlap in time [16–20].

Studies have shown that overlapping compact binary coalescence (CBC) signals can lead to biased results in the measurement of source parameters if not properly accounted for [17–19,21]. In turn, this can negatively impact the inference of CBC mass distributions or tests of general relativity [22]. Therefore, it is important to develop methods to properly account for overlapping signals. A first step in that direction was made in Janquart *et al.* [21], where we have applied joint parameter

estimation (JPE) and hierarchical subtraction (HS) to a population of high-mass binary black hole (BBH) signals. There, we showed that JPE can accurately recover the parameters of the individual signals, while HS works in some cases but often leads to biased results for close merger times. However, the JPE method is computationally expensive, which is a severe inconvenience when accounting for the large detection rate expected in the future. Therefore, we argued that a good strategy would be to apply HS by default and JPE whenever one knows that bias is present if the overlap is neglected. However, *a priori* there is no obvious way of knowing when to expect significant bias. An approach that has been proposed in the literature is to use appropriately adapted Fisher matrices to predict bias that would occur if overlaps are neglected [23]. Here, we build upon this and more recently developed Fisher matrix analyses [24] to also include prior information in the result. We find that the method is reliable for large biases, but that it fails to accurately predict the shift in the posterior probability distributions when they are comparable to the effect of the detector noise. In contrast, we demonstrate that a time-frequency overlap method [25] shows greater reliability in predicting the bias.

Here we propose a different approach where we perform parameter estimation (PE) on a single signal and reweight the posterior to one with the best estimate of the other signal removed from the strain. By calculating the Jensen-Shannon divergence (JSD) between the original and the reweighted posterior we are able to accurately predict if the overlapping signals affected PE. If they did not, we already have the accurate posterior. If they did, we perform JPE to fully account for the overlapping signals.

To make the application of JPE to realistic 3G-like data computationally feasible, we rely on a modification of the PE framework developed in Janquart *et al.* [21], which now includes relative binning (RB) [26–29]. With this, we are able to cope with scenarios including precession, higher-order modes (HOMs), and a lower analysis cutoff frequency of 5 Hz.

This work is structured as follows. In Sec. II we outline the different PE methods and demonstrate JPE with RB produces accurate posteriors. In Sec. III, we show the performance of different methods of estimating bias due to signal overlaps, and we introduce and assess our new PE-based methodology. We close with our conclusions in Sec. IV.

II. PARAMETER ESTIMATION METHODS

The goal of GW parameter estimation for CBCs is to fully characterize the parameters $\vec{\theta}$ of the binary—such as chirp mass $\mathcal{M}_c$, mass ratio q , companion spins a_1, a_2 —by obtaining their posterior probability density function (PDF): $p(\vec{\theta}|d, \mathcal{H}_s)$, conditioned on the observed data d and the hypothesis $\mathcal{H}_s$ that a signal is present in the data and it is accurately described by the model we use. To avoid

clutter, we will drop $\mathcal{H}_s$ from further equations, but it should be understood that every probability is conditioned on it. From Bayes' theorem, we get an expression for the posterior [30]

$$p(\vec{\theta}|d) = \frac{\mathcal{L}(d|\vec{\theta})p(\vec{\theta})}{\mathcal{Z}(d)}, \quad (1)$$

with $\mathcal{L}(d|\vec{\theta})$ being the likelihood of parameters given data, $p(\vec{\theta})$ is the prior probability of parameters, and $\mathcal{Z}(d)$ is the normalization factor called evidence, given by

$$\mathcal{Z}(d) = \int d\vec{\theta} p(d|\vec{\theta})p(\vec{\theta}). \quad (2)$$

For GW parameter estimation, $d(t)$ is the time domain strain data recorded by the interferometer network, split between the signal component $h(t)$ and the noise component $n(t)$

$$d(t) = n(t) + h(t). \quad (3)$$

When we consider overlapping signals, $h(t)$ can be split further between individual signal contributions. In this paper, we limit ourselves to overlaps of two signals, denoted by A and B . With this decomposition, we can rewrite the data strain as

$$d(t) = n(t) + h_A(t) + h_B(t). \quad (4)$$

Assuming colored Gaussian noise, the likelihood takes the form

$$\log \mathcal{L}(d|\vec{\theta}) \propto -\frac{1}{2} \langle d - h(\vec{\theta}) | d - h(\vec{\theta}) \rangle, \quad (5)$$

with the inner product $\langle \cdot | \cdot \rangle$ defined as

$$\langle a | b \rangle = 4\Re \int_{f_{\text{low}}}^{f_{\text{high}}} df \frac{a^*(f)b(f)}{S_n(f)}. \quad (6)$$

Here $a(f)$, $b(f)$ are Fourier transforms of time domain functions $a(t)$, $b(t)$; $*$ denotes the complex conjugate; f_{low} and f_{high} are the limiting frequencies of the analysis, and $S_n(f)$ is the one-sided power spectral density (PSD).

A. Single-signal parameter estimation

The simplest way to handle the actual PE of overlapping signals is to ignore the presence of a second signal, assuming a template for a unique merger as in Eq. (5). Generally, we recover only the parameters of the loudest signal present in the data at that time, possibly with biased estimates due to the other underlying signals. These biases can potentially be sufficiently large that the true parameters are outside of the recovered distribution. If the signals are

close in SNR, we may recover the parameters of the quieter signal instead. For injections performed in this paper, we have never observed the sampler mixing both signals in the posterior. Sampling with Bilby uses differential evolution [31] proposal method for drawing new nested sampling live points. Two overlapping signals have their modes far enough away in the parameter space that the sampler becomes fixated on one of the modes and does not propose jumps between them.

B. Hierarchical subtraction

In hierarchical subtraction, we improve on the previous method by sequentially performing single signal parameter estimation (SSPE) analyses and subtracting best-fitting signals. First, we run the SSPE on the full data containing both signals, $d(A, B)$. We recover the parameters of one of the signals, calling it A (in general, this is the signal with higher SNR): $p(\vec{\theta}_A | d(A, B))$. From the posterior, we can get a point estimate of the parameters of the signal $\hat{\theta}_A$ (such as the maximum likelihood sample), and the associated waveform

$$\hat{h}_A(t) = h(t, \hat{\theta}_A). \quad (7)$$

We can then subtract this waveform from the data to obtain

$$\begin{aligned} d(B, r_A, t) &= d(A, B, t) - \hat{h}_A(t) \\ &= n(t) + h_B(t) + h_A(t) - \hat{h}_A(t) \\ &= n(t) + h_B(t) + r_A(t), \end{aligned} \quad (8)$$

with $r_A(t) = h_A(t) - \hat{h}_A(t)$ the residual due to imperfect recovery of signal A . If the process works correctly, this residual should have an SNR much lower than that of signal B .

We can now run SSPE on the reduced data $d(B, r_A, t)$ to obtain the estimates of the posterior of signal B , $p(\vec{\theta}_B | d(B, r_A))$, and corresponding best point estimate $\hat{\theta}_B$.

We can repeat the steps above, this time subtracting signal B and performing PE on

$$d(A, r_B, t) = d(A, B, t) - \hat{h}_B(t) \quad (9)$$

to get a new estimate of the posterior of the signal A . In principle, we can continue this process iteratively however long we wish.

This approach is simple to implement, as it requires only SSPE runs and it was shown to produce good results in many cases [21]. However, it is not guaranteed to converge to the correct posterior. After subtracting one of the signals, the residual remaining in the strain can still be strong enough to bias PE. Additionally, it cannot account for possible correlations in the posteriors between the signals, as during each PE step the parameters of one of the signals are effectively fixed.

C. Joint parameter estimation

The methods described above all use approximations to simplify the analysis, but we can also model all the signals merging in the data (in the cases considered in this paper there are only 2). We modify the signal template as

$$h(\vec{\theta}) = h_A(\vec{\theta}_A) + h_B(\vec{\theta}_B), \quad (10)$$

with h_A, h_B the individual signal templates. With this, we can modify the likelihood to [21]

$$\begin{aligned} \mathcal{L}(d(A, B) | \vec{\theta}) &\propto \exp \left[-\frac{1}{2} \langle d(A, B) - h_A(\vec{\theta}_A) - h_B(\vec{\theta}_B) | \right. \\ &\quad \left. d(A, B) - h_A(\vec{\theta}_A) - h_B(\vec{\theta}_B) \rangle \right], \end{aligned} \quad (11)$$

with the expanded parameter space $\vec{\theta} = (\vec{\theta}_A, \vec{\theta}_B)$.

Of note is that the above likelihood is symmetric under the exchange of signals A and B , so the posterior will inherit this symmetry. Still, for PE it is useful to break this symmetry between signals completely, to accurately recover the properties of each signal. To this end, we sort the posteriors to enforce an ordering of chirp masses $\mathcal{M}_c$ to follow the ordering in the injection. We check if $\mathcal{M}_{\text{inj},A} > \mathcal{M}_{\text{inj},B}$, then for each sample θ_i we compute $(\mathcal{M}_{i,A} > \mathcal{M}_{i,B}) = (\mathcal{M}_{\text{inj},A} > \mathcal{M}_{\text{inj},B})$. If true, we keep the sample as is; if false, we swap each parameter A with the corresponding parameter B for this sample. This works for signals considered in this work, but for signals with very close chirp masses, some other method of ordering the parameters would need to be employed.

While providing the most accurate method of modeling overlapping signals, this JPE method suffers greatly from computational constraints. The likelihood takes twice as long to compute (as we have to generate two waveforms for each likelihood call), but it is the expansion of parameter space that poses the biggest problem. For standard BBH analysis of a single signal, the parameter space is 15-dimensional. When doubled, the parameter space grows to 30 dimensions which is much higher than usual analysis and slows down parameter estimation considerably. In Janquart *et al.* [21], it took 24 days to complete JPE on a 16-core Intel Skylake Gold 6148 CPU for high-mass overlapping signals and a starting frequency of 20 Hz. With 3G detectors, signals will be longer due to lower starting frequency, and, consequently, PE will be slower. To tackle JPE in a realistic scenario, we need to combine it with a method of speeding up the likelihood evaluation.

D. Relative binning

In general, the frequency domain GW strain in the detector, $h(\theta, f)$ is a highly oscillatory function, which in turn requires high sampling frequency to accurately capture its behavior. However, we do not need to capture

the waveform accurately over all the parameter space $\vec{\theta}$, but just in the region where the likelihood $\mathcal{L}(\vec{\theta}|d)$ is close to its maximum value. Relative binning (RB) is a method of likelihood approximation which exploits this, approximating the ratio of a waveform to a reference waveform by a piecewise linear function [26,28,29]. RB is a form of the heterodyne principle, which was first proposed for use in GW data analysis in Cornish [32].

Let us denote by $\vec{\theta}'$ the reference waveform parameters, and by $h'(f)$ the corresponding reference waveform. $\vec{\theta}'$ is chosen to be close to the maximum of $\mathcal{L}(\vec{\theta}|d)$. While $h(f)$ varies quickly with frequency f of the signal, the ratio $h(f)/h'(f)$ varies much more slowly if we remain close to the maximum likelihood point, requiring fewer frequency points to capture the signal behavior. Let us denote by $b = [f_{\min}, f_{\max}]$ a particular frequency bin in which we evaluate the signal, whose width is chosen such that the waveform ratio can be approximated as

$$\frac{h(f)}{h'(f)} = r_1(b) + r_2(b)(f - f_c(b)) + \mathcal{O}(f^2), \quad (12)$$

and the quadratic correction is negligible. $f_c(b)$ is the central frequency of the bin and $r_i(b)$ coefficients can be calculated from the waveform ratio at the bin edges.

This binning scheme performs well for nonprecessing signals with just the (2,2) mode present [26] and for this case it is implemented in the Bilby package [33]. With higher modes present, the ratio $h(f)/h'(f)$ still varies quickly with frequency and many bins are required to capture the signal. To improve RB, we can apply it to each mode of the signal individually, so that we will not have to worry about capturing the beating pattern between the modes [28,29]. The detector strain can be written as

$$h(f) = \sum_{l,m} (F_+ C_{l,m}^+(f) + F_\times C_{l,m}^\times(f)) h_{l,m}^L(f). \quad (13)$$

Here (l, m) denote component modes of the signal; $(F_+, F_\times)$ are the detector response functions; $C_{l,m}^+(f)$, $C_{l,m}^\times(f)$ are the twisting-up functions, transforming the signal from the coprecessing frame L to the total angular momentum frame; and $h_{l,m}^L(f)$ are the individual modes of the signal in the coprecessing frame. We define $C_{l,m}(f) = F_+ C_{l,m}^+(f) + F_\times C_{l,m}^\times(f)$ and perform RB mode-by-mode

$$\frac{h_{l,m}(f)}{h'_{l,m}(f)} = r_{l,m}^1(b) + r_{l,m}^2(b)(f - f_c(b)) + \mathcal{O}(f^2) \quad (14)$$

$$\frac{C_{l,m}(f)}{C'_{l,m}(f)} = s_{l,m}^1(b) + s_{l,m}^2(b)(f - f_c(b)) + \mathcal{O}(f^2), \quad (15)$$

where $C'_{l,m}(f)$ is the twisting-up function for the reference waveform. $r_{l,m}^i(b)$ and $s_{l,m}^i(b)$ by the ratio of waveform

modes and twisting-up functions at the bin edges. With the above approximation, the inner products in the likelihood function can be approximated by

$$\begin{aligned} \langle d|h \rangle &= \Re \sum_b \sum_{(l,m)} [A_{(l,m)}^0(b) r_{(l,m)}^{0*}(b) s_{(l,m)}^{0*}(b) \\ &\quad + A_{(l,m)}^1(b) (r_{(l,m)}^{1*}(b) s_{(l,m)}^{0*}(b) + r_{(l,m)}^{0*}(b) s_{(l,m)}^{1*}(b))], \\ \langle h|h \rangle &= \Re \sum_b \sum_{(l,m)} \sum_{(\tilde{l}, \tilde{m})} \{ B_{(l,m),(\tilde{l}, \tilde{m})}^0(b) r_{(l,m)}^0 r_{(\tilde{l}, \tilde{m})}^{0*} s_{(l,m)}^0 s_{(\tilde{l}, \tilde{m})}^{0*} \\ &\quad + B_{(l,m),(\tilde{l}, \tilde{m})}^1(b) [s_{(l,m)}^0 s_{(\tilde{l}, \tilde{m})}^{0*} (r_{(l,m)}^0 r_{(\tilde{l}, \tilde{m})}^{1*} + r_{(\tilde{l}, \tilde{m})}^{0*} r_{(l,m)}^1) \\ &\quad + r_{(l,m)}^0 r_{(\tilde{l}, \tilde{m})}^{0*} (s_{(l,m)}^0 s_{(\tilde{l}, \tilde{m})}^{1*} + s_{(\tilde{l}, \tilde{m})}^{0*} s_{(l,m)}^1)] \}, \end{aligned} \quad (16)$$

with coefficients A^0 , A^1 , B^0 , B^1 computed once, before the sampling starts, and defined by

$$\begin{aligned} A_{(l,m)}^0(b) &= \frac{4}{T} \sum_{f \in b} \frac{d(f) h_{(l,m)}^{L*}(f)}{S_n(f)}, \\ A_{(l,m)}^1(b) &= \frac{4}{T} \sum_{f \in b} \frac{d(f) h_{(l,m)}^{L*}(f)}{S_n(f)} (f - f_c(b)), \\ B_{(l,m),(\tilde{l}, \tilde{m})}^0(b) &= \frac{4}{T} \sum_{f \in b} \frac{h_{(l,m)}^L(f) h_{(\tilde{l}, \tilde{m})}^{L*}(f)}{S_n(f)}, \\ B_{(l,m),(\tilde{l}, \tilde{m})}^1(b) &= \frac{4}{T} \sum_{f \in b} \frac{h_{(l,m)}^L(f) h_{(\tilde{l}, \tilde{m})}^{L*}(f)}{S_n(f)} (f - f_c(b)). \end{aligned} \quad (17)$$

We then choose the frequency bins b for our signal by choosing a fiducial waveform different from the reference waveform and ensuring that the difference between the true log-likelihood and the approximate log-likelihood for this waveform falls below the desired error tolerance, 0.01 for our paper. An example of this approach is proposed in Leslie *et al.* [28], where the full frequency range is recursively bisected until the given error tolerance is reached and we use the same approach here.

We use and build upon the computationally efficient implementation in [29,34] where the authors created a fast, Bilby-compatible [35,36] implementation for the IMRPhenomXPHM waveform model [37].

E. Combining joint parameter estimation with relative binning

The advantages of using RB for likelihood evaluations are that it is very fast and, for our use case, that it easily generalizes to multiple signals in the detector. Indeed, the separate spin-weighted harmonics already behave like separate signals, so the only change necessary is to change the indexing in Eq. (16) from (l, m) to (i, l, m) , where i labels the signals present in the detector. We must mention that while this works fine for two signals, it does scale




FIG. 1. RB is applied to two overlapping signals. They were chosen to be nonprecessing and to contain just the dominant (2, 2) mode, for clarity of the figure. Top: frequency domain strain of the individual and the combined signals. They are strongly oscillating, requiring high sampling frequency to be captured properly. Bottom: real part of the ratio of the strain to the reference strain. The ratio of the combined signal to the combined reference signal is much smoother than the absolute signal. The ratios of the composite signals are much smoother and are better suited for approximation by a piecewise linear function than the combined ratio.

quadratically with the number of signals present: the number of waveform inner products we need to calculate to obtain the likelihood scales as $N(N+1)/2$ where N is the number of signals. The number of frequency bins needed to capture the signal stays constant, as we bin each signal individually.

An example of how RB works for multiple signals is shown in Fig. 1. At the top, we plot the real part of the frequency domain strain for the individual and the combined signals. We see that it is strongly oscillating and as such requires high resolution to arrive at a proper representation. In the bottom plot, we show the ratio of the strain to the reference strain. If we consider the ratios of the combined signals, the function is much smoother than the absolute strain and requires fewer points to represent it as a piecewise linear function. But if we choose to represent each signal ratio individually, the functions are even less oscillating, leading to just a few frequency points needed for the likelihood approximation.

For this paper, we modify the mode-by-mode RB implementation from [29] to be compatible with multiple signals and use it throughout to produce our results. This implementation is only compatible with the IMRPhenomXPHM waveform model, as twisting-up procedure is waveform dependent. For completeness we additionally implemented JPE with RB for a generic waveform model without

precession nor higher modes. Details about the code release can be found in Appendix A.

F. Tests of joint parameter estimation with relative binning

For our tests of JPE with RB, we use a detector network comprised of one 10 km triangular ET observatory located at the Virgo site¹ and two CEs located at sites of LIGO-Hanford and LIGO-Livingston having arm lengths of 40 km and 20 km respectively. We inject two BBH signals with the IMRPhenomXPHM waveform model [37] into Gaussian noise following projected noise curves for ET [39] and CE [40]. We set the lower frequency cutoff of the analysis at 5 Hz, for which we note that traditional PE approaches become untractable for large injection studies. We then analyze the signals with the Bilby [35,36] software using Dynesty [41] as nested sampler. We use our custom implementation of the RB likelihood.²

We draw 80 signal pairs from the priors summarized in Table I and inject them into Gaussian noise. We only use injections for which SNRs are above 8. We keep merger times of both signals close to each other to test JPE on cases where overlapping signal bias is expected to be significant [18].

1. Likelihood approximation

We start with assessing how well our RB JPE implementation approximates the true likelihood with the two signals in the strain. For constructing RB frequency bands, we follow [29] for the reference waveform and bin settings, using the injected signal for the former, and a change in 2% for the chirp mass and mass ratio for the latter. We then recursively bisect the full frequency range until the RB error on the likelihood evaluated for the test waveform falls under a predetermined threshold, which we set to 0.03, following the default behavior of the RB code we adapted [34].³ These settings lead to a frequency grid of around a few hundred points, depending on the injected signal. We note that for some injections (around 1 in 5) this leads to a frequency grid that does not capture the likelihood properly, but allowing for a different percentage

¹The design of ET has not been finalized between two L-shaped detectors and one triangular detector, the latter of which we chose for this work. For comparisons between designs, see Branchesi *et al.* [38].

²<https://github.com/corvus314/jointRB>.

³To be more precise, we require that the likelihood error in each bin falls below $0.03/\sqrt{N_b}$, where N_b is the proposed number of bins at each iteration of the bisection algorithm. We run the algorithm until the number of bins and the requested number of bins converge. After our analysis, we discovered that the RB implementation we adapted sometimes stops this process while one more bisection is possible, due to a rounding issue. This has led to excess likelihood error over the desired level. Even with this bug, our likelihoods were accurate enough that none of the posteriors produced in this work were affected.

TABLE I. Summary of generation and analysis priors for tests of JPE. Both overlapping signals share the same priors, with signal *A* chosen to have higher SNR in the pair. The parameters have their usual meanings: chirp mass $\mathcal{M}_c$, mass ratio q , dimensionless spins a_i together with their angles θ_i to the orbital angular momentum, azimuthal angle between them $\Delta\phi$ and the precession angle ϕ_{JL} . RA and DEC stand for the right ascension and the declination while D_L is the luminosity distance to the source. θ_{JN} is the angle between the line of sight and the total angular momentum while ψ is the polarization angle, while ϕ_c and t_c stand for the phase and time at coalescence. The UniformInComponents prior on $\mathcal{M}_c$ and q is equivalent to uniform priors in the component masses m_1, m_2 within the prior boundaries. UniformSourceFrame is the distance prior uniform in comoving volume and the source frame time.

Parameter	Prior
$\mathcal{M}_c$	UniformInComponents ($10M_\odot, 100M_\odot$)
q	UniformInComponents (0.1, 1)
$a_{1,2}$	Uniform (0, 1)
$\theta_{1,2}$	Sin (0, π)
$\Delta\phi, \phi_{JL}$	Uniform (0, 2π)
RA	Uniform (0, 2π)
DEC	Cos ($-\pi/2, \pi/2$)
θ_{JN}	Sin (0, π)
ψ	Uniform(0, π)
ϕ_c	Uniform (0, 2π)
D_L	UniformSourceFrame (0.1 Gpc, 40 Gpc)
t_c	Uniform (-0.05 s, 0.05 s)

difference (typically 10%) in chirp mass and mass ratio fixes the problem.

In Fig. 2 we show how well the RB behaves for an example signal. We plot the RB log-likelihood error $\delta \ln \mathcal{L} = \ln \mathcal{L}_{RB} - \ln \mathcal{L}_{\text{true}}$ against the log-likelihood ratio $\ln \mathcal{L}$ of every sample from the posterior, relative to the maximum likelihood sample $\ln \mathcal{L}_{\text{max}}$. We see that the error on the likelihood clusters around the value 0.325, around 10 times larger than the error tolerance we have chosen. This can be explained by choosing the error tolerance not at any particular point from the posterior (which we do not know before sampling), but an arbitrary point in parameter space close to the injection, which is not the maximum likelihood point, due to presence of noise in the detector.

Even with $\delta \ln \mathcal{L} \approx 0.33$ all the likelihood errors lie close to this value, within the (0.24, 0.41) range. We can show that this is enough to guarantee a good approximation of the likelihood. Starting with samples drawn from the posteriors obtained with the RB likelihood, we can reweight them to the samples drawn from the true unapproximated likelihood by assigning to them weights equal to ratios of their likelihoods. We see from the definition of $\delta \ln \mathcal{L}$ that the expression for the weights is $w_i = \exp(-\delta \ln \mathcal{L}_i)$. From the weights, we can compute the effective sample size [42]




FIG. 2. RB likelihood approximation error for an example overlapping signal. The scatter plot shows the distribution of likelihoods for all the points from the posterior. The x-axis shows the full unapproximated likelihood $\ln \mathcal{L}$ relative to the maximum likelihood sample $\ln \mathcal{L}_{\text{max}}$. The y-axis shows the absolute error $\delta \ln \mathcal{L}$ in the binned log-likelihood over the true likelihood. The relative error between the samples is at most 0.15, which leads to the effective sampling rate of $\eta_{\text{eff}} = 99.96\%$.

$$n_{\text{eff}} = \frac{(\sum_i^N w_i)^2}{\sum_i^N w_i^2}, \quad (18)$$

and the corresponding resampling rate $\eta_{\text{eff}} = n_{\text{eff}}/N$, which measures the fraction of the original N samples we are effectively left with after the reweighting process. For our example signal, this is equal to $\eta_{\text{eff}} = 99.96\%$ —high enough that the posterior obtained with RB is indistinguishable from the true posterior.⁴ For all signals investigated, the resampling rate is above 99.8%, showing us that our implementation accurately approximates the true likelihood of overlapping signals and that any reweighting step is in practice unnecessary.

2. Performance on test injections

With the RB likelihood verified to accurately reflect the true likelihood of the two overlapping signals, we compare the performance of JPE against less computationally expensive methods, to demonstrate that some signals require modeling all the mergers present for accurate inference.

In Fig. 3, we show an example of a posterior with significant bias present. For many event pairs we analyzed, the presence of another signal in the strain has only a small effect on the posteriors. We have deliberately chosen to highlight one where the effect of the overlapping signal is very clear.

JPE posteriors represents what happens when we properly model all the signals present and as such is our ground truth, representing the true shape of the posterior,

⁴This assumes that the RB likelihood effectively covers the true posterior. If there are regions of the posterior for which RB likelihood erroneously sets the likelihood close enough to zero that we do not get any sample from them, η_{eff} would be unaffected. If we instead got even a single sample from those regions, η_{eff} would go to 0.




FIG. 3. Comparison of different PE methods for overlapping signals and a representative signal in Gaussian noise. JPE is compared with SSPE and the second step of the HS method where the dominant signal is analyzed again after subtracting the estimation of signal B from the strain. For comparison, a posterior with only a single signal injected into the strain is also plotted. While both JPE and HS recover similar posteriors, they are statistically distinct. The contours in the corner plot indicate 0.86 credible regions.

with the injected value properly recovered. It is very close to the posterior with no overlapping signal in the strain, a common feature for all 80 injections. In JPE, the uncertainty on the recovered value of θ_B affects the recovery of θ_A , leading to a small difference between the two posteriors.

In contrast, when we do have overlapping signals present and try to recover just one, the resulting recovery tends to be biased. In this example, it is enough to shift the posteriors completely away from the injected value, but smaller shifts are common. In many cases, the biased posterior can be restored closer to its correct position by performing multiple hierarchical subtraction steps (HS II). In this example, subtracting the subdominant signal and reanalyzing the data shifts the posterior close to the JPE posterior. There is still a small difference between the posteriors.

While this could suggest the less computationally expensive hierarchical subtraction could be used over JPE to get the posteriors quicker, we have to be careful as it does not perform well on all the events, like in the example posterior shown in Fig. 4. Here, when we recover overlapping signals with SSPE, the bias is more severe. While hierarchical subtraction reduces it substantially and shifts the posterior closer to JPE, it is not enough and the true values are still completely missed in the recovery. In principle, we could try more hierarchical subtraction steps to further improve the method, but more steps do not make sense from a




FIG. 4. Comparison of different PE methods of overlapping signals for an extreme bias case. JPE is compared with SSPE and the second step of the HS method where the dominant signal is analyzed again after subtracting the estimation of signal B from the strain. For comparison, a posterior with only a single signal injected into the strain is also plotted. While both JPE and HS recover similar posteriors, they are statistically distinct. The contours in the corner plot indicate 0.86 credible regions. For this particular injection, SSPE case has a very small secondary peak closer to the injection, a feature very uncommon in the injections considered in this paper.

computational point of view for two signals. With RB, JPE takes on average 11 days⁵ on a 16-core CPU Intel Skylake Gold 6148, 6 times longer than a SSPE run and 2 times as long as the three PE runs needed for hierarchical subtraction. With two additional steps, most of the computational advantage of hierarchical subtraction would be lost (although others have attempted to speed up HS steps using neural networks [43]). The method may still be useful where there are more than two signals overlapping in the strain.

3. Percentile-percentile test

To assess the accuracy of the overlapping signals posteriors when analyzing them with different PE methods, we perform a percentile-percentile test (pp test) of the posterior consistency [36,44]. We draw injections from the prior, perform parameter estimation with the same prior, get 1D distributions for each of the injection parameters, and record at which percentile in the posterior the injected value was recovered. We then plot the fraction of events that recovered the injected value in the given percentile. If the analysis was unbiased and the posterior samples were accurately drawn from the true posteriors, the percentiles

⁵For our set of injections the average SNR was 53.

should be distributed uniformly and the resulting graph should be a diagonal line (with some deviation due to the finite number of injections).

In Fig. 5, we show how the different methods (all using RB approximation) recover the parameters of the signals. The cited p-values for individual parameters were computed using the Kolmogorov-Smirnov test (KS test) [45] and combined using Fisher’s combined probability test [46].

At the top, we see how badly the standard PE method performed when applied to two overlapping signals. We can only recover the louder signal, resulting in only 80 injections plotted. The pp-plot significantly departs from a straight diagonal line, with a significant fraction of events recovered at 0.0 and 1.0 percentiles. For those events, the bias due to the

presence of the second signal is strong enough that the injection lies completely outside the probability mass of the posterior. This is reflected by the very low p-values for all the parameters, and biases present in all the posteriors. The combined p-value is essentially equal to 0.

In the middle panel, we see how the hierarchical subtraction steps have improved the parameter recovery. After initial SSPE, we subtracted the recovered signal, estimated parameters of quieter signal, and estimated parameters of the louder signal again, with the quieter signal partially removed. This results in 160 total signals plotted. While the pp-plot is somewhat closer to the diagonal, it still shows inconsistent parameter recovery, with some injections recovered completely outside the posterior. The combined p-value is still essentially 0.

At the bottom, we show the result for recovery with JPE, where our data template properly models both mergers present in the strain. As we recover both signals, we plot 160 signals total. We see that the fraction of events in which we completely missed the injection in the recovery properly goes to 0. The p-values are much higher than for the previous methods with the lowest p-value again obtained for a_1 at 0.085. The combined p-value is 0.337, indicating that posteriors obtained from JPE are self-consistent.

We conclude that our extension of RB to overlapping signals works and we can recover accurate posteriors with JPE. Hierarchical subtraction proved unreliable. Posteriors recovered with it are inconsistent, even as they do improve with each additional PE step.

III. ESTIMATION OF BIAS DUE TO OVERLAPPING SIGNALS

Running JPE on all pairs of potentially overlapping signals should lead to accurate posteriors as all the signals would be properly modeled in the likelihood. However, it is computationally inadvisable, as the method is much slower than the standard PE—in this paper we find that JPE on two signals is 6 times slower than SSPE. It is therefore of interest to develop a method that could ascertain if JPE is needed for a given pair of signals.

To test different methods of estimating PE bias due to overlapping signals, we generate an astrophysical population of BBHs and inject the signals into the detector network described in Sec. II F into zero noise, to focus on effects coming from the overlaps. We draw the component masses from the power-law + peak distribution from [6] and the component spins from the DEFAULT spin model from the same paper. We draw the redshifts from the merger rate reconstructed in Oguri’s fit [47]. All the other parameters are drawn from isotropic distributions, except time of coalescence which we set to 0 for all the events.⁶ We draw 1000 superthreshold events ($\text{SNR} > 8$) from our




FIG. 5. Percentile-percentile plot for different analysis methods of overlapping signals. The x-axis is the percentile at which the injection was recovered. The y-axis is the fraction of events for which the injection was recovered for a given percentile. The shaded regions indicate expected deviation from the diagonal at 0.68, 0.95, and 0.997 probability levels. The legend notes the p-values of the individual parameters, while the combined p-value is quoted above the plots. Top: recovery of the signal SSPE template. Only the louder signal can be recovered, resulting in 80 plotted injections. Middle: recovery of both signals using the hierarchical subtraction method, 160 plotted injections. Bottom: recovery of the signal with joint parameter estimation, 160 total signals.

⁶This guarantees a large fraction of events for which overlapping signals affect PE.

population and arrange them into 500 pairs of overlapping signals.

A. Prior-informed Fisher matrices

Fisher-based approaches are common for studies into the effects of overlapping signals in the 3G era [23,25]. We incorporate the prior information into them and compare how well their predictions hold up against the PE posteriors.

The Fisher matrix method relies on applying the linear signal approximation [48] to expand the waveform model $h(\vec{\theta})$ near the maximum likelihood point $\vec{\theta}_{\text{ML}}$

$$h(\vec{\theta}) \approx h(\vec{\theta}_{\text{ML}}) + \partial_i h(\vec{\theta}_{\text{ML}}) \delta\theta^i. \quad (19)$$

For the case of two signals present in the data, with just one modeled by the waveform, the approximated likelihood takes form

$$\log \mathcal{L}(d|\vec{\theta}) \approx -\frac{1}{2}(\theta^i - \theta_{\text{ML}}^i) \Gamma_{ij} (\theta^j - \theta_{\text{ML}}^j), \quad (20)$$

where $\Gamma_{ij} = \langle \partial_i h | \partial_j h \rangle$ is the Fisher information matrix and the maximum likelihood point is located at

$$\theta_{\text{ML}}^i = \theta_A^i + \Gamma_{ij}^{-1} \langle \partial^j h | n \rangle + \Gamma_{ij}^{-1} \langle \partial^j h | h_B \rangle. \quad (21)$$

Here Γ_{ij}^{-1} is the covariance matrix, which is the inverse of the information matrix. We can interpret the second term on the right-hand side as shifting the likelihood peak due to the presence of noise in the data and the third term as shifting the peak due to bias introduced by the presence of the second signal [49]. For the details about the derivation one can refer to Antonelli *et al.* [23], for example.

With the above approximation, we can predict the bias on parameter recovery we would expect from the presence of overlapping signals without performing expensive parameter estimation. According to the expression for the likelihood derived above, we would expect the posterior of parameter θ^i to be a Gaussian centered at θ_{ML}^i with standard deviation $\sigma = \sqrt{\Gamma_{ii}^{-1}}$ (no summation is implied here by the use of repeated indices).

The location of the maximum likelihood point is determined up to the linear order and so breaks down as the bias due to the overlapping signal increases. We account for this by running a differential evolution algorithm [50] to refine our localization of the maximum. We use the Fisher-predicted value as our initial guess and set up a bounding box based on predicted parameter standard deviations. This helps to ensure that the likelihood expansion is performed near the true maximum and so improves the accuracy of the method. We use the implementation from the SciPy package [51].

The standard Fisher-based approach ignores any priors and estimates the posteriors just from the likelihood.




FIG. 6. Estimation of the chirp mass measurements of the dominant signal for different bias estimation methods, as a function of SNR of the subdominant signal. The SNR of the dominant signal is fixed to 80. The error bars and the color band indicate 90% credible interval. The true PE posterior (blue) shifts linearly from the injected value of $\mathcal{M}_{c,A}$ (black) for all values of SNR_B . The standard Fisher method (orange) predicts a posterior that is not shifted enough from the injected value and that is too wide. The prior-informed (PI) Fisher method (green) correctly predicts the posteriors for small SNR_B , but breaks down as the subdominant signal gets louder. Refined PI Fisher method (red) numerically finds the maximum the maximum likelihood point and is similar to the PE posterior even for high SNR.

We follow Dupletsa *et al.* [24] in incorporating the prior information into the Fisher formalism. We sample from the Fisher likelihood truncated by the prior boundaries [52],⁷ and assign to each sample a weight equal to its prior probability $w_i(\vec{\theta}) = p(\vec{\theta})$. The resulting weighted samples behave as if they were drawn from $\mathcal{L}(d|\vec{\theta})p(\vec{\theta})$ distribution instead of $\mathcal{L}(d|\vec{\theta})$ distribution they were drawn from.

We can see the effect of the improvements above in Fig. 6. The standard Fisher method underestimates the shift in the posterior and overestimates its width. Incorporating the prior information makes the Fisher posterior much closer to the PE posterior, as long as the SNR_B of the quieter signal B is not too big. If we refine the location of the likelihood maximum with differential evolution, the Fisher method no longer fails for high SNR_B .

The above comparison was performed for a single draw of the signal parameters. While including the prior information leads to more accurate posteriors estimated with Fisher matrices, it is not a reliable method for all injections. In the following section, we show that the posteriors obtained with this method still differ from the PE posteriors and in some cases do not overlap with them at all.

1. Test of the Fisher matrix method

In this section, we assess our best performing Fisher method—prior-informed Fisher with the maximum likelihood point refined by sampling.

⁷We use the implementation from <https://github.com/brunzema/truncated-mvn-sampler>.




FIG. 7. Comparison of the prior-informed Fisher posteriors (blue) with the PE posteriors (orange) and the injected value (black) for overlapping signals under the hypothesis of one signal present in the interferometer. Top: example event for which the method works well. The significant shift from injection in chirp mass $\mathcal{M}_c$ is well predicted. Middle: example event for which the Fisher method successfully predicts significant shift in sky localization. Bottom: example event for which the Fisher method fails to approximate PE posteriors, which is shifted away from the injection due to overlapping signals.

We can see example 1D posterior for chirp mass ($\mathcal{M}_c$), declination (DEC), and right ascension (RA) in Fig. 7. These posteriors are also representative of other unimodal parameters, except the mass ratio that the Fisher method never recovers accurately. In the top row, we see an example of the prior-informed Fisher matrix method (including numerical estimation of the maximum likelihood point) when it performs well. For most injections, it gives posteriors closely matching PE, even when the posterior is significantly shifted from the injection, like in the case of $\mathcal{M}_c$. In the middle row, we show an example of a posterior where neglecting overlapping signals in the analysis leads to a large error in the sky localization. The Fisher method in general predicts this shift in posterior well, even when its predictions of $\mathcal{M}_c$ start to diverge from PE results. Still, there are cases where the method breaks down, as in the example in the bottom row (we could detect no pattern which parts of the parameter space are prone to this). In particular, the DEC posterior is completely separated from the PE posterior. If we used the Fisher method to predict the bias due to overlapping signals in this case, we would be completely off. We would underestimate the bias due to overlaps and conclude the posteriors are much less affected than they actually are.

To summarize how well the Fisher method performs among the all 500 injections, we quantify bias in each




FIG. 8. Bias $\Sigma_\theta = |\bar{\theta} - \theta_{inj}|/\sigma$ on the parameter estimation due to the presence of another signal, not accounted for in the template. The bias from PE is compared to the bias predicted with Fisher matrix method. The results are shown for chirp mass $\mathcal{M}_c$ and right ascension RA, representative of other parameters. Two different methods of estimating Fisher bias are plotted. One estimates the maximum likelihood point described and the other uses the true maximum likelihood point obtained from PE.

parameter by a single number

$$\Sigma_\theta = |\bar{\theta} - \theta_{inj}|/\sigma, \quad (22)$$

where θ is a parameter of the binary, θ_{inj} is the injected value of the parameter, $\bar{\theta}$ is the mean of the posterior and σ its standard deviation. It quantifies the shift in the posterior relative to its width. We take the absolute value for the convenience of comparison across multiple orders of magnitude. This is the same definition of bias used in Antonelli *et al.* [23] and measures how far the mean of the posterior is located away from the injection. It is not particularly useful for parameters with multiple peaks, so we only use it for parameters with unimodal posteriors.

In Fig. 8 we show how the bias estimated from the Fisher matrix compares to the bias obtained from PE. We chose to highlight the bias on chirp mass $\mathcal{M}_c$ and right ascension, as their behavior is representative of other parameters. We plot the results of Fisher expansion around the maximum likelihood point predicted with Fisher formalism and the location of the true maximum obtained from the sampling. The latter serves as an upper bound on accuracy achievable by Fisher matrix.

Above $\Sigma \approx 1$ we see that the predicted and the true bias cluster close to the diagonal line, indicating they agree rather well. Note however the logarithmic nature of the plot; there can still be a factor of 2 difference between them.

This is due to the sensitivity of our definition of bias Σ_θ to the standard deviation of the posterior for high biases. For example, the posteriors in the upper left corner of Fig. 7 have a difference of $\Sigma_{\mathcal{M}_c}$ of 2.7 even though the predicted absolute shift is the same. Most outlier points in this region get close to the diagonal after using the true maximum likelihood point, indicating that the Fisher approximation used for finding this point starts breaking down for biases far away from the injection.

For $\Sigma \lesssim 1$, so for the bias at or below the effects of the detector noise, the Fisher matrices fail to accurately predict the bias. The Fisher method can predict a relatively substantial bias of $\Sigma \approx 1$ even for small true bias in the 0.001–0.1 range. As such it is not accurate enough to serve as a predictor of whether we need to perform JPE or just a standard PE. The standard deviation of the mean scales as $\sigma/\sqrt{N}$ where σ is the standard deviation of the distribution and N is the number of drawn samples. With N in the 2000–10,000 range, the uncertainty in our definition of bias due to sampling is ~ 0.01 – 0.02 . We would need the Fisher approach to be two orders of magnitude more accurate to predict small biases. These issues persist even when we expand near the true maximum likelihood point so we find Fisher formalism only useful when it predicts biases above the noise level $\Sigma \approx 1$.

B. Time-frequency overlap

Previous work [25] with Fisher matrices indicated that the biases due to overlapping signals occur when time-frequency tracks of the signals intersect, which is equivalent to nonvanishing cross-correlation between the signals. We verify this explicitly with the PE results on our 500 astrophysical injections into zero noise. We plot the Jensen-Shannon divergence (JSD) [53] between posteriors obtained with one and two signals in the strain as a function of the match $|M|$ between the signals A and B

$$M = \langle h_A | h_B \rangle / \sqrt{\langle h_A | h_A \rangle \langle h_B | h_B \rangle}, \quad (23)$$

which is just normalized cross-correlation. JSD is a measure of the statistical distance between distributions p and q . It is defined by

$$\text{JSD}(p, q) = \frac{1}{2} D(p||m) + \frac{1}{2} D(q||m), \quad (24)$$

where $m = (p + q)/2$ is the mixture distribution and

$$D(p||q) = \int dx p(x) \ln \frac{p(x)}{q(x)} \quad (25)$$

is the relative entropy [54] between the distributions. We chose JSD as our distance measure as it is symmetric between the distributions. In Romero-Shaw *et al.* [36] the

authors have shown that when sampling with Bilby package statistical fluctuations of the sampler lead to JSD of up to 0.0015 nats between the marginalized 1D posteriors. We adopt $\text{JSD} > 0.0015$ as our criterion of bias due to presence of overlapping signals.

The results can be seen in Fig. 9. We plot the JSD between the chirp mass posteriors obtained through PE with one and two signals in the strain as a function of the




FIG. 9. JSD between chirp mass posteriors with and without overlapping signals as a function of the match between the signals for different time separation of the mergers $dt = t_B - t_A$ at geocenter. As soon as the signals are separated temporally, the chirp mass ordering starts predicting bias due to overlaps ($\mathcal{M}_{c,A} > \mathcal{M}_{c,B}$ mergers are expected to overlap in time-frequency for $dt > 0$). The horizontal line at 0.0015 nats indicates JSD expected due to statistical fluctuations of the sampler.

match between the signals. We see that more similar signals (higher match) lead to larger distortion of the posteriors due to overlaps (higher JSD) as predicted.

We also plot the same results where we shifted the separation of the signals and kept all the other parameters the same, and we see an interesting pattern emerge. At $dt = 0$ s both signals are coincident at geocenter but it does not mean that they are coincident at the detectors. Depending on the sky positions of the signals, the ordering of the signals at the detectors can be different at each of the detectors. But as soon as the time difference between the mergers is larger than 0.045 s (the Earth's diameter in seconds) the ordering of the signal at the interferometers is consistent between the interferometers and independent of sky location. The match and JSD are then driven primarily by ordering of the chirp masses of the signals, consistent with biases occurring due to overlaps in time-frequency. The chirp mass determines the leading order frequency evolution of the signal [55]

$$\dot{f} = \frac{96}{5} \pi^{8/3} \mathcal{M}_c^{5/3} f^{11/3}. \quad (26)$$

Up to leading order, the conditions $t_B > t_A$ and $\mathcal{M}_{c,A} < \mathcal{M}_{c,B}$ imply that the time-frequency tracks never cross, resulting in small match and small effect of overlaps on PE.

In Fig. 10 we demonstrate that the chirp mass ordering can act as a fairly accurate estimator of bias due to overlapping signals even for close mergers (as long as they are further apart than 0.045 s). For $dt = 0.05$ s the chirp mass ordering leads to 86% accurate classification of merger pairs with no bias and 88% accurate classification of merger pairs with bias.




FIG. 10. Percentage of injections for which presence of overlapping signals leaves chirp mass posteriors unbiased ($JSD < 0.0015$ compared with single injected signal posterior). For $dt > 0.045$ s chirp mass ordering can act as a fairly accurate indicator of bias.

C. Reweighting single-signal posteriors

Here we propose a novel method to determine biases due to overlapping signals based on the initial PE performed on the data. It does not rely on Fisher matrices and thus takes the full shape of the likelihood into account. In this paper, we focus on pairs of overlapping BBH signals, but the method is easily generalizable.

We start by going over the list of triggers from the search pipeline and identifying all the pairs of signals overlapping in time.⁸ Search pipelines will still be able to perform this task in the 3G era [56]. We choose the higher SNR merger in the pair and perform the standard PE algorithm, completely ignoring the presence of overlapping signals—we use the likelihood $\mathcal{L}(\vec{\theta}_A|d)$. We then subtract the highest likelihood estimate of the signal $h(\vec{\theta}_A)$ from the strain. For the other event, we use the differential evolution algorithm [50] to find the maximum of $\mathcal{L}(\vec{\theta}_B|d - h(\vec{\theta}_A))$ and set it as our estimate of its parameters $\vec{\theta}_B$. The time needed to find the approximate maximum is negligible compared to the time needed for the sampling. We then use this estimate to reweight our posterior of signal A to one with signal B removed from the strain. Explicitly, we assign to each sample a weight

$$w_i = \frac{\mathcal{L}(\vec{\theta}_A^i|d - h(\vec{\theta}_B^i))}{\mathcal{L}(\vec{\theta}_A^i|d)}. \quad (27)$$

If the reweighted posterior is the same as the original posterior ($JSD < 0.0015$ nats for all 1D posteriors), the presence of overlapping signals was not important for inference. We can perform standard PE on signal B and we spend the same amount of computational resources as if there was no possibility of overlaps. If the posteriors are different, we run the full JPE analysis on the pair. Compared with starting from JPE from the beginning, computational cost increases around 17% by requiring the initial PE run. This does not pose a problem, as only a small fraction of mergers in the 3G era are expected to be affected by overlapping signals. It is also much lower than around 500% increase in the computational cost if we run JPE when only standard PE run is needed. In Fig. 11, we show a schematic flowchart of our proposed method.

1. Test of the posterior reweighting method

In Fig. 12 we show how the JSD predicted by our method ($JSD_{\max}$) compares to the true JSD between the single-signal posterior and the overlapping signals posteriors ($JSD_{\max}$) based on our 500 injections into zero noise. We see that reweighting PE posteriors leads to an accurate

⁸For the general deployment of this method. For the purpose of this paper, we assume all our injections triggered the search pipelines.




FIG. 11. Schematic flowchart on how to determine if overlapping signals caused bias in PE.

estimate of JSD between them, as long as we use the 1D posteriors of all our sampling parameters. It is less reliable if we use just a single parameter due to low reweighting efficiency in cases of significant bias. When the effects of overlaps shift the posteriors significantly in one parameter (high true JSD), the original posterior poorly covers the target posterior resulting in a low effective sample size. This leads to high predicted JSD for all parameters, even the ones that should be unaffected by the overlaps. If we use maximum predicted JSD from across all the parameters, this high predicted JSD will only happen when the true JSD is high as well and high precision is not important. This effect is responsible for higher scatter around the line of equal JSD for signal pairs with high JSD.

We see an interesting effect in areas of low true JSD, below the 0.0015 nats expected from statistical fluctuations during the sampling. The JSD is consistently underestimated in this region. This shows us that for these mergers, the effect of accounting for overlaps is orders of magnitude lower than fluctuations due to sampling. When we reweight




FIG. 12. The predicted JSD ($\widehat{JSD}_{\max}$) between overlapping and nonoverlapping signal posteriors plotted against the true JSD ($JSD_{\max}$) for 500 injections into zero noise. For a single parameter (top) it follows a clear trend line. As we take maximum JSD from more parameters (middle), the prediction gets better until we add all the sampling parameters (bottom). The points are colored according to the resampling rate η_{eff} to the reweighted posterior. The blue lines indicate a JSD of 0.0015 nats.

the posterior we are relying on the same realization of the sampler and do not get up to 0.0015 nats of JSD difference between the distributions we get from our PE run that is used for calculating the true JSD.

An important part of each figure is the lower-right quadrant created by the lines of $JSD = 0.0015$ nats, which in the classification task represents wrongly concluding the overlaps cause no bias in PE when they do distort the shape of the posterior. There are only 5 points in this region indicating $\sim 1\%$ false-negative rate for our method. For a science-conservative approach, we want the false-negative rate to be as close to zero as possible. To that end, we could move our JSD threshold to a lower number, like 0.0004 nats. This would be enough to accurately classify all but one cases. For a computation-conservative approach, we may decide that the 1% false-negative rate is acceptable, especially as the true JSD of the points in this region is still low,




FIG. 13. The predicted JSD ($\widehat{JSD}_{\max}$) between overlapping and nonoverlapping signal posteriors plotted against true JSD ($JSD_{\max}$) for 500 injections into Gaussian noise. For a single parameter (top) it follows a clear trend line. As we take maximum JSD from more parameters (middle), the prediction gets better until we add all the sampling parameters (bottom). The points are colored according to the resampling rate η_{eff} to the reweighted posterior. The blue lines indicate a JSD of 0.0015 nats.

below 0.003 nats. This would lead to a minimal bias in the recovered posteriors, see Fig. 14.

The above discussion was based on our injection into zero-noise. In Fig. 13 we repeat the process for the same injections, but injected into Gaussian noise. It follows a similar pattern to the zero noise injections with a wider scatter around the line of accurately predicted JSD. As such, the accuracy of the method drops and the false-negative rate rises to 2%. Again, this can be lowered by changing our JSD cutoff at the cost of higher false-positive rate. In Appendix B, we show how different implementation choices affect the performance of our method.

IV. CONCLUSIONS

Mergers of binary black holes with close coalescence times in the 3G era can cause significant biases in parameter estimation if the overlapping signals are unaccounted for.

If the bias is present, hierarchical subtraction can shift the posteriors closer to their true value in most cases, but the resulting posteriors are still distorted, with HS failing the pp-test. We need to properly model all signals with joint parameter estimation to recover accurate posteriors.

We adapted joint parameter estimation to work with relative binning for IMRPhenomXPHM as well as waveform models without higher modes nor precession. For binary black holes, it produces reliable parameter estimation results in the timeframe of 11 days for a pair of overlapping signals on a 16-core Intel Skylake Gold 6148 CPU (for average signal SNR of 53). This can be used in future injection studies into Einstein’s Telescope and Cosmic Explorer to properly model the effects of overlaps on the recovered properties of the signals.

We have tested how reliable prior-informed Fisher matrices are for the estimation of bias due to overlaps and found them reliable only when they predict bias above the value expected from the detector noise. The overlap of time-frequency tracks on the other hand proved to be a good indicator of the presence of bias due to overlaps. Even for mergers with $dt = 0.05$ s it accurately predicted if bias was present in 86% of cases.

Finally, we proposed a parameter estimation-based approach to determining if the presence of overlapping signals causes bias. By reweighting the posterior to one with other signals subtracted and quantifying the shift, we can determine if the bias is present with 99% accuracy in zero noise (98% accuracy in Gaussian noise). If one is willing to pay the price of a few false alarms, one can push this to 100% identification of cases needing joint parameter estimation. The choice depends on whether one prefers a science-conservative approach, or a computing-conservative approach. This method does not rely on any simplifications to the likelihood inherent in the Fisher formalism and properly takes into account the realization of the strain in the detector, but is much slower, being PE-based. However, compared with the resources that are anyway needed to run parameter estimation on the signals, it does not introduce significant overhead.

In future work, our method will be tested on signals with higher temporal separations. Due to the computational constraints of running parameter estimation, we have focused on binary black hole mergers close together in time. This was to guarantee a population for which bias due to overlaps played a significant role so that different parameter estimation techniques could be tested, and our method of predicting bias had a wide range of biases to verify it. We leave it for future work to test the method on binary neutron stars, overlaps of which are expected to be common in the 3G era.

The plots were prepared with Matplotlib [57]. NumPy [58] and SciPy [51] were used during data analysis.

ACKNOWLEDGMENTS

We thank Qian Hu for useful comments on the manuscript. T.B., H.N., A.S., and C.V.D.B. are supported by the research program of the Netherlands Organisation for Scientific Research (NWO). T.D. acknowledges funding from the EU Horizon under ERC Starting Grant, No. SMART-101076369. The authors are grateful for computational resources provided by Cardiff University, and funded by STFC Grants No. ST/I006285/1 and No. ST/V005618/1 supporting UK Involvement in the Operation of Advanced LIGO.

Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX A: CODE RELEASE

We made our Fisher matrix and Joint Parameter Estimation codes available at [github](https://github.com/corvus314/jointRB)⁹ together with examples showing their use. They were developed to work together with Bilby [35,36] parameter estimation software.

While our JPE code was extensively tested on the case of 2 signals, we developed it to work with an arbitrary number of signals present. We do note that this may cause a noticeable decrease in performance.

JPE with higher modes requires reimplementing some functions from LALSimulation [59] and as such our implementation is hardcoded for the IMRPhenomXPHM waveform [37]. We also include a separate implementation for a general LALSimulation waveform model, as long as it does not include higher modes or precession. While not a focus of this paper, this second implementation can give a quick parameter estimation of overlapping binary neutron stars.

APPENDIX B: ADDITIONAL DETAILS ON BIAS ESTIMATION METHOD

In this paper, we chose to quantify the bias due to overlapping signals with JSD between single-signal and overlapping posteriors. Compared with other measures, such as the shift of the parameter average from the injection, it is not an intuitive measure. As such in Fig. 14 we show examples of posterior pairs with different levels of JSD between them.

⁹<https://github.com/corvus314/jointRB>.




FIG. 14. Example pairs of chirp mass posteriors with different levels of JSD between them.

In our method, we rely on performing PE just on a single signal and finding the best estimate of the maximum likelihood parameters of the other signal. This mismatch between the true and the estimated signal has only small effects on the performance of our method. We can see this in Fig. 15, where there is almost no difference between the two.




FIG. 15. Performance of our bias estimation method when using the true parameters of signal B compared to using estimates of the maximum likelihood parameters. There is almost no difference in the posterior.




FIG. 16. Adaptation of the bias estimation method for case where SSPE was performed on both overlapping signals. Top: subdominant signal, is estimated by its maximum likelihood sample. Middle: reweighting is averaged over 100 samples with the highest likelihood from SSPE of the subdominant signal. Bottom: reweighting is averaged over 100 draws from the SSPE posterior of the subdominant signal.

If we were willing to perform initial PE on both signals instead of just one, our method could be improved further. Instead of reweighting the posterior of event A to one estimate of signal B, we could draw multiple samples, reweight the likelihood to each one, and then average it over the samples. If we are willing to pay the initial PE cost, this

approach is still fast and parallelizable. We show the results in Fig. 16 based on taking 100 samples from the posterior of signal B. If we take the samples at random from the posterior, the method performs worse than with just the maximum likelihood estimate. If we take 100 samples with the highest likelihood, the method shows minor improvements.

-
- [1] J. Aasi *et al.* (LIGO Scientific Collaboration), Advanced LIGO, *Classical Quantum Gravity* **32**, 074001 (2015).
 - [2] F. Acernese *et al.* (Virgo Collaboration), Advanced Virgo: A second-generation interferometric gravitational wave detector, *Classical Quantum Gravity* **32**, 024001 (2015).
 - [3] The LIGO Scientific, the Virgo, and KAGRA Collaborations, LIGO/Virgo/KAGRA public alerts, GraceDB (2025).
 - [4] R. Abbott *et al.* (KAGRA, Virgo, and LIGO Scientific Collaborations), GWTC-3: Compact binary coalescences

- observed by LIGO and Virgo during the second part of the third observing run, *Phys. Rev. X* **13**, 041039 (2023).
- [5] R. Abbott *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), Tests of general relativity with GWTC-3, [arXiv:2112.06861](https://arxiv.org/abs/2112.06861).
- [6] R. Abbott *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), The population of merging compact binaries inferred using gravitational waves through GWTC-3, *Phys. Rev. X* **13**, 011048 (2023).

- [7] R. Abbott *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), Constraints on the cosmic expansion history from GWTC-3, *Astrophys. J.* **949**, 76 (2023).
- [8] Report of the LSC Post-O5 study group, <https://dcc.ligo.org/LIGO-T2200287/public> (accessed: 2025-03-20).
- [9] Kentaro Somiya (KAGRA Collaboration), Detector configuration of KAGRA: The Japanese cryogenic gravitational-wave detector, *Classical Quantum Gravity* **29**, 124007 (2012).
- [10] Yoichi Aso, Yuta Michimura, Kentaro Somiya, Masaki Ando, Osamu Miyakawa, Takanori Sekiguchi, Daisuke Tatsumi, and Hiroaki Yamamoto (KAGRA Collaboration), Interferometer design of the KAGRA gravitational wave detector, *Phys. Rev. D* **88**, 043007 (2013).
- [11] Bala Iyer *et al.*, LIGO-India, Proposal of the Consortium for Indian Initiative in Gravitational-Wave Observations (IndIGO), <https://dcc.ligo.org/LIGO-M1100296/public> (2021).
- [12] M. Punturo *et al.*, The Einstein Telescope: A third-generation gravitational wave observatory, *Classical Quantum Gravity* **27**, 194002 (2010).
- [13] Adrian Abac *et al.*, The science of the Einstein Telescope, [arXiv:2503.12263](https://arxiv.org/abs/2503.12263).
- [14] David Reitze *et al.*, Cosmic Explorer: The U.S. Contribution to gravitational-wave astronomy beyond LIGO, *Bull. Am. Astron. Soc.* **51**, 035 (2019), <https://baas.aas.org/pub/2020n7i035/release/1>.
- [15] B. Sathyaprakash *et al.*, Scientific objectives of Einstein Telescope, *Classical Quantum Gravity* **29**, 124013 (2012).
- [16] Tania Regimbau and Scott A. Hughes, Gravitational-wave confusion background from cosmological compact binaries: Implications for future terrestrial detectors, *Phys. Rev. D* **79**, 062002 (2009).
- [17] A. Samajdar, J. Janquart, C. Van Den Broeck, and T. Dietrich, Biases in parameter estimation from overlapping gravitational-wave signals in the third-generation detector era, *Phys. Rev. D* **104**, 044003 (2021).
- [18] E. Pizzati, S. Sachdev, A. Gupta, and B. S. Sathyaprakash *et al.*, Toward inference of overlapping gravitational-wave signals, *Phys. Rev. D* **105**, 104016 (2022).
- [19] P. Reilton and V. Raymond, Parameter estimation bias from overlapping binary black hole events in second generation interferometers, *Phys. Rev. D* **104**, 084039 (2021).
- [20] Yoshiaki Himemoto, Atsushi Nishizawa, and Atsushi Taruya, Impacts of overlapping gravitational-wave signals on the parameter estimation: Toward the search for cosmological backgrounds, *Phys. Rev. D* **104**, 044010 (2021).
- [21] J. Janquart, T. Baka, A. Samajdar, T. Dietrich, and C. Van Den Broeck, Analyses of overlapping gravitational wave signals using hierarchical subtraction and joint parameter estimation, *Mon. Not. R. Astron. Soc.* **523**, 1699 (2023).
- [22] Qian Hu and John Veitch, Accumulating errors in tests of general relativity with gravitational waves: Overlapping signals and inaccurate waveforms, *Astrophys. J.* **945**, 103 (2023).
- [23] A. Antonelli, O. Burke, and J. R. Gair, Noisy neighbours: Inference biases from overlapping gravitational-wave signals, *Mon. Not. R. Astron. Soc.* **507**, 5069 (2021).
- [24] Ulyana Dupletsa, Jan Harms, Ken K. Y. Ng, Jacopo Tissino, Filippo Santoliquido, and Andrea Cozzumbo, Validating prior-informed Fisher-matrix analyses against GWTC data, *Phys. Rev. D* **111**, 024036 (2025).
- [25] Aaron D. Johnson, Katerina Chatzioannou, and Will M. Farr, Source confusion from neutron star binaries in ground-based gravitational wave detectors is minimal, *Phys. Rev. D* **109**, 084015 (2024).
- [26] B. Zackay, L. Dai, and T. Venumadhav, Relative binning and fast likelihood evaluation for gravitational wave parameter estimation, [arXiv:1806.08792](https://arxiv.org/abs/1806.08792).
- [27] L. Dai, T. Venumadhav, and B. Zackay, Parameter estimation for GW170817 using relative binning, [arXiv:1806.08793](https://arxiv.org/abs/1806.08793).
- [28] N. Leslie, L. Dai, and G. Pratten, Mode-by-mode relative binning: Fast likelihood estimation for gravitational waveforms with spin-orbit precession and multiple harmonics, *Phys. Rev. D* **104**, 123030 (2021).
- [29] H. Narola, J. Janquart, Q. Meijer, K. Haris, and C. Van Den Broeck, Relative binning for complete gravitational-wave parameter estimation with higher-order modes and precession, and applications to lensing and third-generation detectors, *Phys. Rev. D* **110**, 084085 (2024).
- [30] J. Veitch and A. Vecchio, Bayesian coherent analysis of inspiral gravitational wave signals with a detector network, *Phys. Rev. D* **81**, 062003 (2010).
- [31] Cajo J. F. Ter Braak, A Markov Chain Monte Carlo version of the genetic algorithm differential evolution: Easy Bayesian computing for real parameter spaces, *Stat. Comput.* **16**, 239 (2006).
- [32] Neil J. Cornish, Fast Fisher matrices and lazy likelihoods, [arXiv:1007.4820](https://arxiv.org/abs/1007.4820).
- [33] Kruthi Krishna *et al.*, Accelerated parameter estimation in Bilby with relative binning, [arXiv:2312.06009](https://arxiv.org/abs/2312.06009).
- [34] J. Janquart and H. Narola, Relativebilbying: A package for relative binning with Bilby, <https://github.com/lemnisl2/relativebilbying> (2022).
- [35] G. Ashton *et al.*, Bilby: A user-friendly Bayesian inference library for gravitational-wave astronomy, *Astrophys. J. Suppl. Ser.* **241**, 27 (2019).
- [36] I. M. Romero-Shaw *et al.*, Bayesian inference for compact binary coalescences with Bilby: Validation and application to the first LIGO—Virgo gravitational-wave transient catalogue, *Mon. Not. R. Astron. Soc.* **499**, 3295 (2020).
- [37] Geraint Pratten *et al.*, Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes, *Phys. Rev. D* **103**, 104056 (2021).
- [38] Marica Branchesi *et al.*, Science with the Einstein Telescope: A comparison of different designs, *J. Cosmol. Astropart. Phys.* **07** (2023) 068.
- [39] S. Hild *et al.*, Sensitivity studies for third-generation gravitational wave observatories, *Classical Quantum Gravity* **28**, 094013 (2011).
- [40] Benjamin P. Abbott *et al.* (LIGO Scientific Collaborations), Exploring the sensitivity of next generation gravitational wave detectors, *Classical Quantum Gravity* **34**, 044001 (2017).
- [41] Joshua S. Speagle, Dynesty: A dynamic nested sampling package for estimating Bayesian posteriors and evidences, *Mon. Not. R. Astron. Soc.* **493**, 3132 (2019).
- [42] Victor Elvira, Luca Martino, and Christian P. Robert, Rethinking the effective sample size, *Int. Stat. Rev.* **90**, 525 (2022).

- [43] Qian Hu, Hierarchical subtraction with neural density estimators as a general solution to overlapping gravitational wave signals, [arXiv:2507.05209](#).
- [44] J. Veitch *et al.*, Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library, *Phys. Rev. D* **91**, 042003 (2015).
- [45] Frank J. Massey, The Kolmogorov-Smirnov test for goodness of fit, *J. Am. Stat. Assoc.* **46**, 68 (1951).
- [46] R. A. Fisher, *Statistical Methods for Research Workers* (Oliver & Boyd, Edinburgh, 1925).
- [47] Masamune Oguri, Effect of gravitational lensing on the distribution of gravitational waves from distant binary black hole mergers, *Mon. Not. R. Astron. Soc.* **480**, 3842 (2018).
- [48] Lee S. Finn, Detection, measurement and gravitational radiation, *Phys. Rev. D* **46**, 5236 (1992).
- [49] Curt Cutler and Michele Vallisneri, LISA detections of massive black hole inspirals: Parameter extraction errors due to inaccurate template waveforms, *Phys. Rev. D* **76**, 104018 (2007).
- [50] R. Storn and K. Price, Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces, *J. Global Optim.* **11**, 341 (1997).
- [51] P. Virtanen *et al.*, SciPy 1.0: Fundamental algorithms for scientific computing in Python, *Nat. Methods* **17**, 261 (2020).
- [52] Z. I. Botev, The normal law under linear restrictions: Simulation and estimation via minimax tilting, *J. R. Stat. Soc. Ser. B* **79**, 125 (2016).
- [53] J. Lin, Divergence measures based on the Shannon entropy, *IEEE Trans. Inf. Theory* **37**, 145 (1991).
- [54] S. Kullback and R. A. Leibler, On information and sufficiency, *Ann. Math. Stat.* **22**, 79 (1951).
- [55] B. P. Abbott *et al.*, The basic physics of the binary black hole merger GW150914, *Ann. Phys. (Berlin)* **529**, 1600209 (2017).
- [56] P. Relton, A. Virtuoso, S. Bini, V. Raymond, I. Harry, M. Drago, C. Lazzaro, A. Miani, and S. Tiwari, Addressing the challenges of detecting time-overlapping compact binary coalescences, *Phys. Rev. D* **106**, 104045 (2022).
- [57] J. D. Hunter, Matplotlib: A 2D graphics environment, *Comput. Sci. Eng.* **9**, 90 (2007).
- [58] Charles R. Harris *et al.*, Array programming with NumPy, *Nature (London)* **585**, 357 (2020).
- [59] LIGO Scientific, Virgo, and KAGRA Collaborations, LVK Algorithm Library—LALSuite, Free software (GPL) (2018).