

Sequential simulation-based inference for gravitational wave signals

Uddipta Bhardwaj^{1,*}, James Alvey^{2,†}, Benjamin Kurt Miller^{2,3,4}, Samaya Nissanke^{1,5} and Christoph Weniger²

¹GRAPPA Institute, Anton Pannekoek Institute for Astronomy and Institute of High-Energy Physics, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands

²GRAPPA Institute, Institute for Theoretical Physics Amsterdam, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands

³AMLab, Institute for Informatics, LAB42, Science Park 900, 1098 XH Amsterdam

⁴AI4Science, Institute for Informatics, LAB42, Science Park 900, 1098 XH Amsterdam

⁵Nikhef, Science Park 105, 1098 XG Amsterdam, The Netherlands



(Received 19 April 2023; accepted 3 August 2023; published 25 August 2023)

The current and upcoming generations of gravitational wave experiments represent an exciting step forward in terms of detector sensitivity and performance. For example, key upgrades at the LIGO, Virgo and KAGRA facilities will see the next observing run (O4) probe a spatial volume around four times larger than the previous run (O3), and design implementations for, e.g., the Einstein Telescope, Cosmic Explorer, and LISA experiments are taking shape to explore a wider frequency range and probe cosmic distances. In this context, however, a number of very real data analysis problems face the gravitational wave community. For example, it will be critical to develop tools and strategies to analyze (among other scenarios) signals that arrive coincidentally in detectors, longer signals that are in the presence of nonstationary noise or other shorter transients, as well as noisy, potentially correlated, coherent stochastic backgrounds. With these challenges in mind, we develop PEREGRINE, a new sequential simulation-based inference approach designed to study broad classes of gravitational wave signal. In this work, we describe the method and implementation, before demonstrating its accuracy and robustness through direct comparison with established likelihood-based methods. Specifically, we show that we are able to fully reconstruct the posterior distributions for every parameter of a spinning, precessing compact binary coalescence using one of the most physically detailed and computationally expensive waveform approximants (SEOBNRv4PHM). Crucially, we are able to do this using only 2% of the waveform evaluations that are required in, e.g., nested sampling approaches. Finally, we provide some outlook as to how this level of simulation efficiency and flexibility in the statistical analysis could allow PEREGRINE to tackle these current and future gravitational wave data analysis problems.

DOI: [10.1103/PhysRevD.108.042004](https://doi.org/10.1103/PhysRevD.108.042004)

I. INTRODUCTION

Observational status. The gravitational wave (GW) sky is now much louder than it was after the first detection of a gravitational wave signal from the merger of two black holes in 2015 [1]. Indeed, source catalogs are now complete and varied enough to be used for studying questions in gravitational theory [2], cosmology [3–5], and even the formation and internal properties of black holes and neutron stars [6,7]. To give some current context, the LIGO-Virgo collaboration (LVKC) has now reported on 90 compact binary coalescence events¹ that are believed to be of astrophysical origin [12,13].

Along with more efficient search pipelines [14–16], one of the major reasons for the increase in size of the gravitational wave catalog is the large increase in search volume probed by the detectors. For example, in the previous LVKC observing run (O3), the projected comoving search volume for binary black hole (BBH) mergers was $V_{\text{BBH}} = 3.4 \times 10^8 \text{ Mpc}^3$ [17]. In comparison, however, the next observing run (O4) that is planned for May 2023 has a projected search volume of around $V_{\text{BBH}} = 1.5 \text{ Gpc}^3$ [17,18], representing over a 400% increase in volume and the corresponding expected event rate.² This observing run is predicted to have the advanced LIGO (aLIGO) detectors running close to their design sensitivities (O5), and therefore, represents a key benchmark for the current generation of gravitational wave observatories.

*u.bhardwaj@uva.nl

†j.b.alvey@uva.nl

¹Additional catalogs such as the deep extended catalog [8], the 3-OGC catalog [9] and others [10,11] report several other sources with different analysis pipelines or higher false alarm rate thresholds.

²The equivalent volumes for binary neutron star (BNS) coalescence events are $V_{\text{BNS}} = 3.3 \times 10^6 \text{ Mpc}^3$ and $V_{\text{BNS}} = 3.4 \times 10^8 \text{ Mpc}^3$ for O3 and O4 respectively [17,18].

Such an improvement in detector sensitivity and corresponding event rate induces a new era of computational challenges when it comes to gravitational wave data analysis.

Gravitational wave data analysis In general, the data analysis pipeline for gravitational waves splits into two parts: detection and parameter inference. While studying the optimality of search pipelines [14,15], especially for new classes of signal, is an interesting and important problem, in this work we will focus on the follow up step of high precision parameter inference. The traditional (and current) approach to gravitational wave parameter inference follows a Bayesian framework where the goal is to infer either the full (N -dimensional) joint posterior over all parameters or marginal posterior distributions over a (sub)set of relevant intrinsic (e.g., source properties) and extrinsic parameters (e.g., the source-detector geometry) [19,20]. In addition, one must make some informed choice (ideally using some astrophysically motivated range, or making a maximally uninformative choice in the absence of this) of the prior distributions for the parameters of interest. For current state-of-the-art waveform approximants/generators, which are now capable of describing complex scenarios such as fully precessing and spinning compact binary inspirals [21–24], the number of relevant parameters is around $N_{\text{params}} \sim 12\text{--}17$. At a purely practical level, this means that approaches such as profiling the gravitational wave data likelihood are computationally infeasible, and it is common therefore to use some form of stochastic sampling method to investigate the parameter space.

Traditional stochastic algorithms such as Markov Chain Monte Carlo (MCMC) [25,26] and nested sampling [27–30] are both implemented within the standard open source analysis software such as LALInference [31], PyCBC [32], and bilby [33–36]. Both MCMC and nested sampling algorithms function by generating samples from the full joint posterior (although they achieve this in different ways), and then typically marginalize to obtain posteriors over individual, or pairs of parameters. The computational cost in these cases comes about since these algorithms carry out numerous likelihood evaluations for each individual sample point. Indeed, a general rule of thumb for current gravitational wave parameter inference is that generating around 10^4 independent posterior samples requires at least 10^6 likelihood evaluations³ (and corresponding forward evaluations of the waveform generation model) [35,36]. Depending on the signal duration and complexity of the waveform approximant, this can result in inference times of hours to weeks for a single events, see e.g. [35,36]. This is the context in which concrete data analysis challenges will arise as, among other things, (i) the detection rate increases significantly [17], (ii) more realistic waveform approximants that incorporate, e.g., higher order

modes become more expensive to evaluate [23,37], and (iii) future gravitational wave detections involving multiple overlapping waveforms [38–42], nonstationary noise distributions, or some stochastic background, see, e.g., [43–45], will render explicit likelihood evaluations increasingly complicated and costly.

Simulation-based inference Recently, there have been significant advances in the field of simulation-based inference (SBI) [46–48] (sometimes also known as implicit likelihood inference), partly as a result of rapid developments in machine learning, but also as a response to emerging data analysis challenges such as those described above. The field of simulation-based inference is now wide and varied, with numerous approaches available as fully implemented open-source software [49–52]. It has been shown to be successful in a number of different physics settings, such as CMB analyses [53], strong lensing image analysis [54], point source searches [55], and field-level cosmology [56], along with other examples [57–61]. Broadly, the key advantages of this set of methods are (i) they can be highly simulation-efficient compared to traditional methods, and (ii) they are “simulation-based” in the sense that they do not require an explicit likelihood to be written down, only a realistic forward simulator to be provided. Combined, this opens up the exciting possibility of using higher fidelity forward models to take full advantage of even the most complex data. In the context of gravitational waves, a particular version of simulation-based inference known as neural posterior estimation (NPE) has already been successfully applied to perform fully amortized inference⁴ on compact binary mergers [62,63]. This leads to impressive inference performance across, e.g., the LVKC catalog, where after the complex training process (requiring millions of waveform evaluations) has been carried out, almost real-time inference can be done [64]. It is worth noting in this regard that the focus of our analysis pipeline is different. We look to perform analysis on individual gravitational wave signals, where this sort of global amortization becomes infeasible. The prototypical example of this would be the overlapping waveform scenario [38–42], where the number of waveform calls required to perform reliable amortized inference would likely be at least an order of magnitude larger, if not more. In this current work, we however focus on single BBH mergers to develop our method because this is where clear, robust quantitative comparisons can be performed to validate our approach.

Key contributions. In this work, we develop a new data analysis method that is applicable to wide classes of gravitational wave signals. In particular, we implement a particular, sequential simulation-based inference algorithm

³This is especially true for spinning and precessing compact binary inspirals.

⁴Amortized inference refers to training a neural network to approximate the posterior distribution over model parameters which allows for fast inference on any input data.

known as truncated marginal neural ratio estimation (TMNRE) [50] based on the SWYFT software [49]. We demonstrate the applicability, accuracy and robustness of this approach by studying two case studies—a highly spinning, precessing binary black hole system at distances of 200 Mpc (extremely high signal-to-noise ratio (SNR) of ~ 100 to act as a stringent test of the precision we can achieve) and 900 Mpc (low SNR of ~ 20 to test our ability to reliably distinguish between signal and noise). Using a state-of-the-art waveform approximant SEOBNRv4PHM [21] (although in practice, any waveform approximant for, e.g., IMRPhenomXPHM [65], IMRPhenomPV2, or IMRPhenomD [22,66], can easily be used), we demonstrate that we can fully reproduce the posterior distributions generated by current likelihood-based methods. Importantly, we show that we can achieve this agreement *with only 2% of the waveform evaluations* that are needed in standard likelihood-based approaches. As far as the computational efficiency of the method is concerned, we provide a discussion of the design choices relevant for our TMNRE-based algorithm in Sec. II, in particular focusing on the details of the network training and optimization. We highlight the role that truncation plays in achieving efficient inference in Fig. 3, where we use an active learning approach to lower the training data variance with respect to the observation of interest. We also discuss the robustness tests that can be carried out to validate our results independent of any other method. In addition, we will release our implementation, which we call PEREGRINE, as open-source, extendable analysis software available to the community.

Structure of the work. The remainder of the work is structured as follows: we begin in Sec. II by briefly reviewing simulation-based inference approaches before giving a detailed explanation of TMNRE and its application to gravitational wave analysis. We apply this method in Sec. III to the two case studies detailed above. Finally, in Sec. IV, we provide some discussion and outline our conclusions regarding the application of our pipeline to current and future gravitational wave data analysis challenges.

II. SIMULATION-BASED INFERENCE FOR GRAVITATIONAL WAVES

In this section, we will give a very brief overview of the general class of simulation-based inference methods, before focusing our attention on the particular implementation we use in this work. Finally, we will discuss some specific design choices that are relevant to the analysis of gravitational waves.

A. Overview of simulation-based inference

Over the last three to four years, there has been a significant increase in the usage and development of so-called simulated-based inference approaches to data

analysis [46,47].⁵ While there are now multiple options in terms of particular implementations and algorithms, every approach has a common setup. In particular, the main question driving all SBI algorithms is: *can we still do robust Bayesian inference if we are only given some generative model?* [46]. In other words, given some forward model $p(x, \theta)$ that takes some underlying set of parameters θ and produces some simulated data x , can we construct meaningful posteriors $p(\theta|x)$ for all, or some subset, of the parameters θ . In practice, how this is done is by studying Bayes' theorem applied to the posterior,

$$p(\theta|x) = \frac{p(x|\theta)p(\theta)}{p(x)}, \quad (1)$$

where $p(\theta|x)$ is the posterior of θ given some observed or simulated data x , $p(x|\theta)$ is the likelihood of a given set of data x given some input parameters θ , $p(\theta)$ is the (Bayesian) prior over the parameters θ , and $p(x)$ is the Bayesian evidence. The key functionality of simulation-based inference is to realize that having a forward generative model $p(x, \theta) = p(x|\theta)p(\theta)$ is equivalent to being able to sample from the (simulated) likelihood. This is the origin of the terms “likelihood-free” or “implicit likelihood” inference [46].

At this point, the various SBI methods diverge somewhat in terms of how they use this capability of sampling from the likelihood to construct posterior densities or posterior samples. Broadly, they can be categorized in terms of how they split up Bayes' theorem in Eq. (1). There are three main ways to do this given a prior $p(\theta)$ (which is explicitly known/chosen) and a generative model $p(x, \theta) = p(x|\theta)p(\theta)$.

- (i) *Neural posterior estimation (NPE)*. The first class of methods attempts to directly estimate the posterior density $p(\theta|x)$, typically through the use of some flexible density estimator such as a normalizing flow [51,62,69,70]. This has been used in a number of contexts, including for the amortized analysis of gravitational waves [63,64,71].
- (ii) *Neural likelihood estimation (NLE)*. In the second case, given some training data generated from the forward model, one attempts to create a well-controlled proxy for the simulated data likelihood $p(x|\theta)$ [52,61,69]. This can then be used directly in traditional algorithms such as MCMC to generate posterior samples.

⁵As discussed in the introduction, this is arguably for two main reasons, the first is that the huge computational improvements in machine learning (ML) methods allow for density estimators, classifiers and neural networks to be efficiently trained on very high dimensional data. The second reason is then that utilizing these developments opens up the possibility to tackle new data analysis challenges [38,67,68].

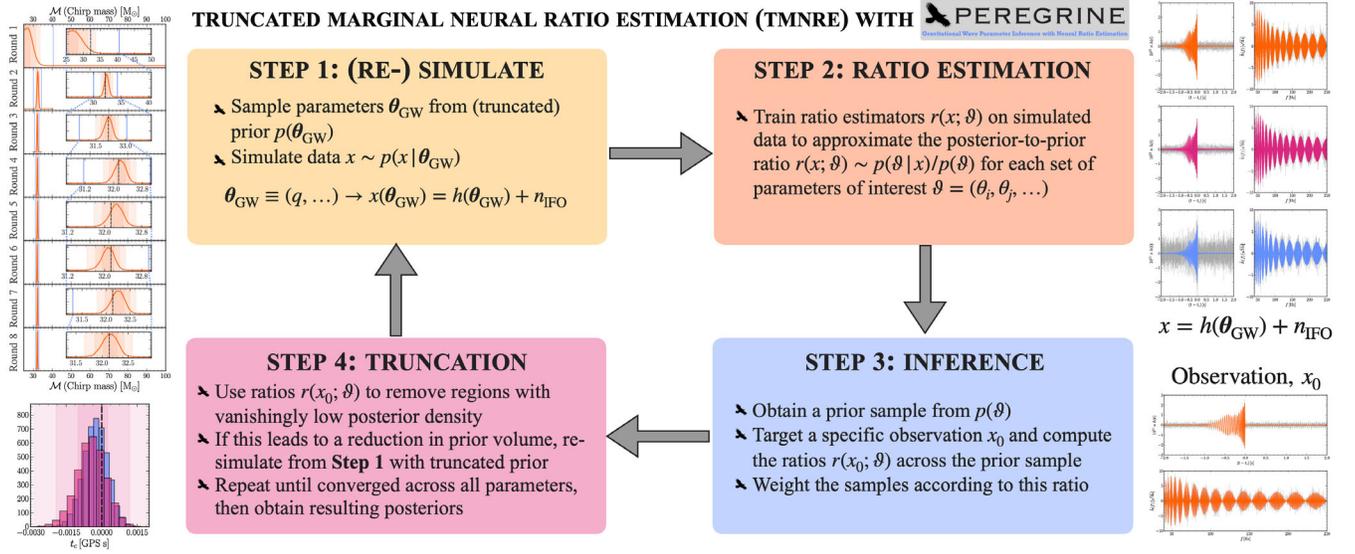


FIG. 1. A schematic illustration of the data analysis method developed in this paper. We use truncated marginal neural ratio estimation (TMNRE) to perform full parameter estimation on gravitational wave signals. The algorithm is implemented in our new code PEREGRINE.

(iii) *Neural ratio estimation (NRE)*. The final class of methods takes a different approach and estimates the likelihood-to-evidence ratio $p(x|\theta)/p(x)$ [49,53–55,57,58,60,72–76]. In this work, we will be using a particular version of this known as truncated marginal neural ratio estimation (TMNRE) [50]. Ratio estimation ultimately works by translating the Bayesian inference problem into a binary classification task on joint and marginal samples. We provide more details of this mapping below in the specific context of gravitational waves.

B. The TMNRE algorithm

In this work, we use a particular simulation-based inference algorithm known as TMNRE [50], implemented within the SWYFT software package [49]. There are two key features of the method that we look to take advantage of for the highly simulation efficient analysis of gravitational wave signals. For reference, the overall structure of the algorithm is described in Fig. 1.

(i) *The “T”*. TMNRE is a sequential method in the sense that the analysis is performed in rounds (typically around 5 to 10). At the end of each round, estimates of the posterior are generated for a *specific*, targeted observation x_0 . The priors are then truncated before generating training data for the next round, leading to a significant (empirically, of the order of 50% to 90%) reduction in simulation budget [53]. This, of course, has the consequence that we are generally not developing an amortized method, but rather one that can be applied to individual observations.

(ii) *The “M”*. The marginal aspect of TMNRE relates to the fact that the algorithm estimates *marginal posteriors*. In other words, instead of estimating the full (N -dimensional) joint posterior $p(\theta|x)$ for all parameters $\theta \equiv (\theta_1, \dots, \theta_N)$, and then marginalizing, we directly estimate the marginal posterior $p(\theta_i, \theta_j, \dots|x)$ for some parameter (or set of parameters) of interest $\{\theta_i, \theta_j, \dots\}$. This again has the advantage that it typically reduces the simulation budget by a large fraction and results in much simpler training procedures because the data-parameter manifolds for the low dimensional marginals are much less complex.

More technically, TMNRE sets up the estimation of the ratio $r(x; \vartheta = (\theta_i, \theta_j, \dots)) = p(x|\vartheta)/p(x)$, where ϑ represents a single parameter or some subset of parameters, as a binary classification problem through the following observation about Bayes’ theorem,

$$r(x; \vartheta) = \frac{p(x|\vartheta)}{p(x)} = \frac{p(\vartheta|x)}{p(\vartheta)} = \frac{p(x, \vartheta)}{p(x)p(\vartheta)}. \quad (2)$$

In other words, we can view this likelihood-to-evidence ratio in multiple ways, (i) as the posterior-to-prior ratio, which will ultimately allow us to generate posterior samples by generating samples from the prior $p(\vartheta)$ and weighting them by the ratio $r(x; \vartheta)$, and (ii) as the ratio between the joint probability density $p(x, \vartheta)$ and marginal density $p(x)p(\vartheta)$.

It is the last form that allows us to carry out a binary classification task. In particular, suppose we have N samples from our simulator [or joint distribution

$p(x, \vartheta) = p(x|\vartheta)p(\vartheta)$ $\{(\vartheta = (\theta_i, \theta_j, \dots), x)_i\}_{k=1\dots N}$, then we can also construct samples from the marginal distribution $p(x)p(\vartheta)$ by randomly shuffling the pair components. The method then attempts to optimize a classifier $d_\phi(x, \vartheta)$ with some trainable parameters ϕ to output a class $d_\phi(x, \vartheta) = 0$ (say) when (x, ϑ) is drawn marginally, and $d_\phi(x, \vartheta) = 1$ when (x, ϑ) is drawn jointly.

In terms of TMNRE, this is technically achieved by minimizing the (binary cross-entropy) loss function [49,73,74],

$$\mathcal{L}[f_\phi] = - \int dx d\vartheta [p(x, \vartheta) \ln(\sigma(f_\phi(x, \vartheta))) + p(x)p(\vartheta) \ln(1 - \sigma(f_\phi(x, \vartheta)))] \quad (3)$$

where $\sigma(x) = [1 + \exp(-x)]^{-1}$ is the sigmoid function and $d_\phi(x, \vartheta) = \sigma(f_\phi(x, \vartheta))$ in the notation above. Crucially, one of the key motivations for TMNRE is that one can actually minimize this loss analytically to obtain the optimal classifier $f_\phi^*(x, \vartheta) \equiv \ln r(x; \vartheta)$ [49]. This means that by successfully optimizing this loss function, *we can directly obtain full information about the likelihood-to-evidence ratio or posterior-to-prior ratio.*

In practice, this is where the ‘‘neural’’ aspect of TMNRE comes into play. Since the data x (and parameters ϑ) can in principle be extremely high-dimensional, we require both very flexible classifiers $d_\phi(x, \vartheta)$ that can process complex data structures such as images or long time series, and also an efficient way to optimize the classifier parameters ϕ (here neural network weights). Recent advances in machine learning architectures, training, and hardware capabilities have now made this process extremely scalable, and we make use of them here, as implemented in the SWYFT package [49,50].

It is important to understand how we can achieve marginal posterior estimates with this algorithm, since it will be important for solving future sampling challenges such as those associated to the analysis of multiple waveforms [38–42]. In TMNRE, we implicitly fully marginalize over every parameter $\theta = (\theta_1, \theta_2, \dots)$ by varying them all in each simulation. When we train the classifier(s), however, we only show the network the small subset of parameters we are interested in, i.e., $(x, \vartheta = (\theta_i, \theta_j, \dots))$. When one analytically minimizes the loss, one will find that you automatically obtain the properly marginalized posterior $p(\vartheta|x)$. In practice, in this work, we will only be interested in obtaining 1- and 2-dimensional marginal posteriors, i.e., $\vartheta \equiv \theta_k$ for some k or $\vartheta \equiv (\theta_i, \theta_j)$ for some pair (i, j) .

We can finally turn to the inference and truncation part of TMNRE. With a fully trained ratio estimator $\hat{r}(x; \vartheta) \simeq p(\vartheta|x)/p(\vartheta)$, we can now target a specific observation x_0 (coming from, e.g., an LVKC data stream, or some mock test observation). By generating samples from the prior $p(\vartheta)$, we can obtain marginal posterior inference results by

weighting these samples with the ratio estimator evaluated on the target observation $\hat{r}(x_0; \vartheta) \simeq p(\vartheta|x_0)/p(\vartheta)$. Inevitably, this will lead to regions of parameter space where the posterior density is high (covering the true values if the algorithm is run successfully) and regions where the density is vanishingly low. We use these regions of low density, specifically where the posterior estimate drops below some threshold $\hat{p}(\vartheta|x_0)/\max_\vartheta \hat{p}(\vartheta|x_0) < \epsilon$ to truncate our prior⁶ before generating samples for the next round of training and inference⁷ [50]. After a number of rounds, once the sensitivity is at the level of statistical uncertainty, this truncation procedure will converge, and the algorithm terminates. This process is highlighted below for the specific case study presented in this work in Fig. 3.

To summarize, the TMNRE algorithm splits into four steps:

- (1) *Step 1:* We generate a batch of simulations $(\vartheta, x)_{i=1\dots N}$ according to our current proposal prior $p(\vartheta)$ using our forward simulator.⁸
- (2) *Step 2:* We train classifier(s)⁹ $f_\phi(x, \vartheta)$ to tell the difference between joint and marginal samples for the parameter(s) of interest ϑ . When optimized, this gives a direct estimate of the ratio $r(x; \vartheta) = p(\vartheta|x)/p(\vartheta)$ for the parameters of interest $\vartheta = (\theta_i, \theta_j, \dots)$.
- (3) *Step 3:* We use the trained ratio estimator to perform inference on a specific target observation x_0 and obtain marginal ratio estimates $\hat{r}(x_0; \vartheta)$.
- (4) *Step 4:* By evaluating the ratios on a set of prior samples, we can truncate the initial proposal distribution to exclude regions of vanishingly low posterior density such that the variance in the training data for the next round is significantly lower.
- (5) We then go back to *Step 1* and repeat until the truncation regions and posterior estimates stabilize, after which we can perform the final inference on the observation x_0 .

This whole process is illustrated diagrammatically in Fig. 1.

⁶This introduces a certain error into the marginal posterior estimation proportional to the integrated quantity $\int_{-\Gamma} d\vartheta p(\vartheta|x_0)$, where $-\Gamma$ is the area outside the new proposal and $\bar{\vartheta}$ is the set of parameters θ not including the parameter(s) being inferred. But, it is specifically in this region that the posterior density is vanishingly low, so the error induced is small and controlled by ϵ .

⁷We typically take $\epsilon \sim 10^{-5}$ which approximately corresponds to a $\sim 4.8\sigma$ exclusion if the posterior was a pure Gaussian. We checked that the inference results are not sensitive to this choice provided ϵ is small enough, only affecting the rate at which the truncation occurs.

⁸Note that this step can be fully parallelized (and indeed is in our implementation), and therefore made arbitrarily fast given hardware access.

⁹In this work, we simultaneously train one classifier for each gravitational wave parameter θ_i , or pair (θ_i, θ_j) , so we end up with a set of classifiers.

TABLE I. Definitions and prior choices for all relevant gravitational wave parameters θ_{GW} in this work. [†] Note that these are the default priors used in BILBY analyses for BBH systems (subject to calibration ranges of the waveform approximants). * Specifically, the luminosity distance prior is taken to be uniform in comoving volume in the *source frame*.

Parameter	Prior choice	Injection
Mass ratio, q	U(0.125, 1)	0.8858
Chirp mass $\mathcal{M}[M_{\odot}]$	U(25, 100)	32.14
Inclination angle θ_{jn} [rad]	sine(0, π) [†]	0.4432
Polarization angle ψ [rad]	U(0, π)	1.100
Phase ϕ_c [rad]	U(0, 2π)	5.089
Tilt angles θ_1, θ_2 [rad]	sine(0, π) [†]	1.497, 1.102
Dimensionless spins a_1, a_2	U(0.05, 1)	0.9702, 0.8118
Spin angles ϕ_{12}, ϕ_{jl} [rad]	U(0, 2π)	6.220, 1.885
Right ascension α [rad]	U(0, 2π)	5.556
Declination δ [rad]	cosine($-\pi/2, \pi/2$) [†]	0.071
Merger time t_c [GPS s]	U(-0.1, 0.1)	0.000
Luminosity Distance d_L [Mpc]	U _{vol} (100, 2000)*	900.0 (C1), 200.0 (C2)

C. Design choices for gravitational waves

There are a number of concrete design choices and implementation steps that we must make in order to apply the TMNRE algorithm to the analysis of gravitational waves. Specifically we must (i) choose/build the forward simulator that generates the data x given parameters θ and some noise model, (ii) implement a neural network that can parse and process the output of the simulator for training the classifier, (iii) choose the relevant prior distributions for our parameters, and (iv) make choices for the relevant parameters in the TMNRE algorithm.

Forward simulator. In terms of the generative model, in this work we choose to study gravitational wave signals from compact binary black hole mergers observed in LIGO and Virgo detectors. We analyze the signal as it appears in all three detectors simultaneously,¹⁰ across both the time and frequency domains. To model the waveforms from the merger, we use the functionality in the open-source code BILBY [33–35]. This allows us to inject a gravitational wave signal $h(\theta_{\text{GW}})$ —typically split into the + and \times -polarization strains and then projected—into a noisy detector with colored noise strain n_{IFO} generated from a power spectral density (PSD) $S_n(f)$. Here, $\theta_{\text{GW}} \equiv (q, \mathcal{M}, \dots)$ is some set of parameters sampled from the prior $p(\theta_{\text{GW}})$ that describe the intrinsic properties of the source under consideration such as the mass ratio q , (redshifted) chirp mass of the system \mathcal{M} , as well as extrinsic properties including the angular position on the sky (α, δ) and the luminosity distance d_L . All of these parameters are provided as input into the waveform

approximant of choice.¹¹ We provide a full list of physical parameters in Table I, including their prior choices. The full forward model for a single detector is then defined by:

$$p(x, \theta_{\text{GW}}) = p(x = h + n_{\text{IFO}} | h_{+, \times}, \theta_{\text{GW}}) \times p(h_{+, \times} | \theta_{\text{GW}}) p(\theta_{\text{GW}}), \quad (4)$$

where h here is the projection of the $h_{+, \times}$ strains onto the detector frame according to the extrinsic parameters in θ_{GW} . The full implementation of the forward model can be found in the `Simulator` class of the PEREGRINE package. Two example observations $x_0 = h(\theta_{\text{GW}}) + n_{\text{IFO}}$ which are the direct output of the simulator and will form the basis of our case studies below are shown in Fig. 2.

Inference Network. To carry out the ratio estimation step in the TMNRE algorithm (Step 2 above), we need to construct a flexible enough parametrization of the classifier $f_{\phi}(x, \vartheta)$. In principle (at least in some infinite training data limit), any sufficiently general parametrization should be able to minimize the loss in Eq. (3). In practice, making sensible design choices about network architectures that are well matched to the data format (particularly the signal structure) leads to huge increases in performance, robustness and general applicability of the method.

In the context of gravitational waves, the signal is a complex time series or frequency domain strain—e.g., at a sampling frequency of 2048 Hz a 4 s signal will consist of over 8000 data points. Most of the information about the

¹⁰Specifically, the Livingston (L1), Hanford (H1) and Virgo (V1) detectors, although any combination or subset can be analyzed also.

¹¹Here we use the SEOBNRv4PHM waveform approximant [21] so as to provide case studies on the most complex spinning and precessing BBH merger systems. Importantly, the analysis pipeline remains identical regardless of this choice, although the specific parameters in θ_{GW} may change.

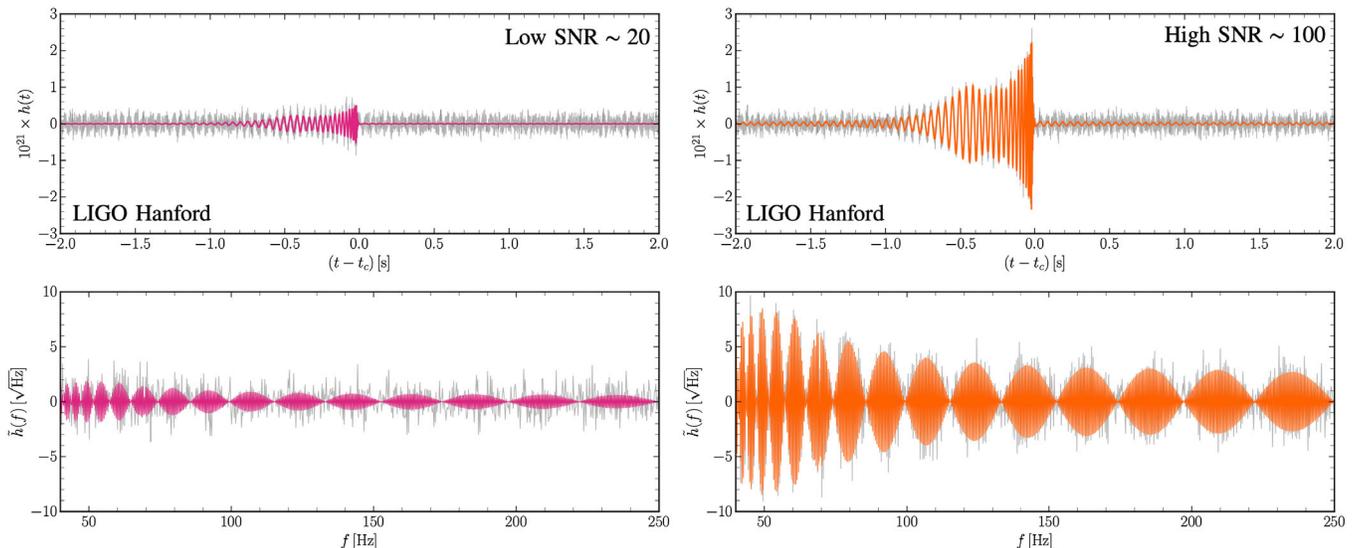


FIG. 2. Two example observations that will be used in the case studies below. The parameter values for the observations can be found in the inference results below (see Figs. 4 and 5) Left: a (typical) low signal-to-noise ratio (SNR) signal as seen in the LIGO Hanford detector. The signal is shown in pink, while the full detector response including noise is shown in gray. Right: a similar signal (shown in orange), but with a very high SNR, again in the Hanford detector. Upper and lower panels: the upper panels illustrate the time-domain strain $h(t)$, while the lower panels show the corresponding frequency domain response $\tilde{h}(f)$, both of which are used in our analysis.

parameters θ_{GW} is encoded in various nonlinear ways. As such, there is no good analytic method for fully processing noisy data streams, or deriving an optimal architecture. On the other hand, in the context of TMNRE, it is empirically the case that the precise network architecture is not so important provided it is (a) sensibly adjusted to the data format—e.g., it would be nonoptimal to first flatten an image, totally removing all spatial structure—and (b) deep enough (i.e., with enough network parameters or structure) to be able to flexibly fit the parameter and data manifold. This relative simplicity comes about as a result of the specific targeting both of single observations and marginal posteriors, significantly reducing the complexity of the fit we are implicitly trying to perform. An interesting comparison is with the NPE implementation of gravitational wave inference where a highly optimized and bespoke network architecture is required to perform fully amortized inference on the high dimensional joint posterior [64].

Practically, we choose a network architecture that acknowledges a number of things about the physics behind gravitational wave signals:

- (i) *Signal segments.* Broadly, compact binary coalescence gravitational wave signals can be split into 3 components (both in the time and frequency domains) corresponding to different phases of their evolution: the inspiral, merger and ringdown phases. The various parameters in θ_{GW} affect these components differently, for example, the coalescence time t_c defines the time of merger, typically associated to the peak of the amplitude in the time domain signal $h(t)$. We want to choose a network that does not treat each part of the waveform identically—such as a

convolutional structure that simply applies the same kernel to each part. Instead we construct a 1d-analog of the well-known unet architecture [77], which is excellent for image (or signal here) segmentation and subsequent analysis on the various segments.

- (ii) *Domain choice.* Information about gravitational wave parameters is encoded differently in the time and frequency domains, although they are related directly by a Fourier transform. At the end of the day, we are interested in extracting some optimal set of features from the signal that are correlated with various physical parameters. Naturally, the effect of some parameters is much more direct in one domain vs. the other. For example, it is very simple to see the effect of varying the time of coalescence on the time domain strain $h(t)$ —it is simply a horizontal shift and is easy to pick up with any simple network. On the other hand, in the frequency domain, the waveform gets modified as $\tilde{h}(f; t_c) = \tilde{h}(f; 0) \exp(2\pi i f t_c)$, which affects the full structure of the frequency domain strain and is much more complicated to extract. In practice we take advantage of the full flexibility of simulation-based methods and present *both* the time and frequency domain strains to the network, analyzing both of them independently to extract features before combining these summary statistics to perform ratio estimation.
- (iii) *Multiple detectors.* The current generation of detectors is set up to provide complementary information about the same gravitational wave event at each of the relevant facilities (e.g., LIGO Hanford and

Virgo). We model the detector response in all (or a subset) of these locations and stack the time (and corresponding frequency) domain strains such that the same geocentric times (or corresponding frequency bins) are aligned. Again, this is taking advantage of the freedom to use any data representation that fully encodes our knowledge about the relevant timing delays in the detectors.

The above describes the rationale for our choice of inference network used to process our data. Its general structure consists of two parallel `unet` networks which act on the time and frequency domain strains independently, before compressing down to a smaller set of summary features in each case. These features, which importantly are *automatically* optimized by the classifier during training are then concatenated and passed to a simple multilayer perceptron that is the default implemented in `SWYFT` which performs the ratio estimation. More generally, one can understand the network as first performing some optimized data compression into a set of summary statistics¹² on which the ratio estimation is then performed with (a subset of) the physical parameters $\vartheta \subset \theta_{\text{GW}}$. The full details of the network architecture can be found both in Appendix A as well as in the `PEREGRINE` implementation. In terms of generalization, we expect this network to perform well on any compact merger signal, including, e.g., overlapping waveforms, although in principle it can also be modified arbitrarily to search for other classes of gravitational waves without any other change in the data analysis pipeline.

Prior choices. The aim of this work is to establish and develop the TMNRE algorithm in the context of gravitational wave analysis. As such, we choose priors $p(\theta_{\text{GW}})$ on the physical parameters $\theta_{\text{GW}} = (q, \mathcal{M}, \dots)$ that represent the standard choices in the literature for analyzing this type of BBH merger signal [20,31–33]. This allows for a direct comparison to inference results from traditional algorithms that are the focus of the next section. For reference, these can be found in, e.g., the `BILBY` documentation [28,30,33–35]. All the parameter definitions and prior distributions relevant to this work can be found in Tab. I.

TMNRE setup. Finally, there are a number of choices that need to be made when running the TMNRE algorithm. These choices are fully detailed in Table II, or in the configuration files supplied with `PEREGRINE`. The settings broadly split into two categories: those that control the training of the inference network, and choices for the TMNRE algorithm itself. In the former case, we can define how many epochs to train the network for (min./max. training epochs), how many epochs to wait for the

TABLE II. Choices for the hyperparameters and settings for the TMNRE algorithm in this work. * This is the minimum number of rounds, if the algorithm has not converged, we continue rounds of inference until the truncation procedure terminates.

TMNRE setting	Value
Number of rounds	7*
Simulation schedule	30k, 60k, 90k, 120k, 120k, 150k, 150k
Bounds threshold ϵ	10^{-5}
Noise shuffling	True
Min./Max. training epochs	30/200
Early stopping	7
Initial learning rate	5×10^{-4}
Training/Validation batch size	256/256
Train: Validation ratio	0.9: 0.1

validation loss to decrease¹³ (early stopping), the initial learning rate (initial learning rate), the batch sizes shown to the network (training/validation batch size), and finally the split in the data between training and validation (train: validation ratio). In addition, we have the TMNRE settings such as the (minimum) number of rounds (number of rounds), the number of simulations to use in each round¹⁴ (simulation schedule), and the bounds threshold for performing the truncation (ϵ). The last setting to discuss is the noise shuffling switch. This is closely related to the question of overfitting to training data, where for small simulation batches it can be the case that the network essentially “remembers” the training data including the noise realizations. This leads to a very low value of the loss on the training dataset, but poor performance on validation data. We found that one extremely effective way of overcoming this is to shuffle the noise realizations in each batch. This effectively shows the network brand new examples every epoch, but with the same signal component relevant for inference. This strategy should be applicable for *any* additive noise model (related to gravitational waves, or not) and can reduce the simulation budget by potentially an order of magnitude.

To summarize this section, we have discussed broadly the field and applicability of simulation-based inference techniques, before describing in detail the development of our new gravitational wave data analysis pipeline `PEREGRINE`. In doing so, we explained the technicalities of our particular

¹²There are some clear analogues with approximate Bayesian computation here [78], although the summary statistics are trained to be optimal rather than hand crafted.

¹³Specifically, the training keeps track of the loss on the training and validation sets, then tests whether the validation loss has increased relative to the previous round. This can be a sign of overfitting, which should of course be avoided, but we are also training via stochastic gradient descent so we instead look for a number of epochs (early stopping) where no decrease in the validation loss has been seen before stopping. We then reinitialize the network to the state with the lowest validation loss.

¹⁴In later rounds, it is typically the case that a larger simulation budget is required to obtain the necessary precision in the parameter inference.

implementation of simulation-based inference, TMNRE, that is applicable to targeted parameter inference on a specific observation of interest. In the next section, we will demonstrate the application of our analysis method to two case studies of highly spinning, precessing BBH mergers, which represent a current state-of-the-art analysis and modeling challenge.

III. RESULTS: TWO CASE STUDIES

We now turn our attention to the application of the method to two concrete case studies that represent the state-of-the-art for parameter estimation in current generation gravitational wave detectors. Specifically, in this section we consider two noisy mock observations with waveforms generated by the SEOBNRv4PHM approximant [21]. This time-domain waveform approximant is tuned to describe BBH merger systems that are both spinning and precessing. It employs more accurate aligned-spin as well as precessing-spin two-body dynamics than its predecessors.¹⁵ It takes as input 15 relevant physical parameters, which are detailed in Table I. We do note, however, that our pipeline makes no assumption about the choice of waveform approximant, so any other model could be easily integrated.

We generate the noise according to the LIGO-O4 power spectral density for each detector (H1, L1, V1) [35]. We assume the detector noise to be Gaussian-distributed, stationary and colored by the estimated PSD of the upcoming LIGO-Virgo observing run.¹⁶ In practice, particularly because we are looking to analyze individual observations, the noise model could be modified to include, e.g., any online estimate of the PSD at the time of an event [1,80]. While of course, the precise posteriors obtained would be different as a result of the different noise realization, the data analysis pipeline would remain identical and we expect our approach to generalize to any noise model.

In the rest of the section we will describe the case studies in more detail before presenting our results when analyzing the signals using the TMNRE algorithm implemented in PEREGRINE. Importantly, we show that we can achieve excellent agreement compared to traditional established methods across all parameters in the 1-dimensional and 2-dimensional marginals. To obtain this comparison, we run the nested sampling code DYNESTY [30] that comes as

¹⁵In particular, it includes the $(2, \pm 2)$, $(2, \pm 1)$, $(3, \pm 3)$, $(4, \pm 4)$ and $(5, \pm 5)$ modes in the co-precessing frame of the binary instead of the commonly used $(2, \pm 2)$ and $(2, \pm 1)$ modes, and displays improved accuracy in calibration fits with numerical relativity simulations across a larger parameter space. When generating aligned-spin waveforms, SEOBNRv4PHM reduces to SEOBNRv4HM.

¹⁶In full technicality, the detector noise originates from a number of different continuous and transient sources which lead it to be nonstationary and colored. However, it has been shown that away from these transient noise artefacts a Gaussian distribution colored by an online estimate of the PSD is an excellent approximation [79].

default with the BILBY package¹⁷ [33–35]. We also discuss additional consistency checks that we can perform on simulation-based inference methods that are independent of—and typically also infeasible to—classical methods.

A. Case study description

We test our data analysis methodology on two separate case studies, both within the context of the SEOBNRv4PHM waveform approximant [21], which is currently one of the most detailed and computationally intensive models available (although, again, we could use any model, see, e.g., [22,23]). In particular we choose an injected signal with the true values given in Table I, which are chosen because they are one set of best fit parameters for the interesting source recently discussed in Refs. [82–84]. In particular, the source comprises two objects that appear to be both precessing and highly spinning (with dimensionless spins $a_i > 0.8$).¹⁸ We place the systems at two different luminosity distances to set up our two case studies:

- (i) C1. $d_L = 900$ Mpc—we study a quieter source (with an SNR of ~ 20) to ensure that our method can deal reliably with inference on a data stream where there is significant amounts of noise. In other words, we test the ability of our method to distinguish between signal and noise in a gravitational wave event.
- (ii) C2. $d_L = 200$ Mpc—conversely, we analyze a very loud source (with an SNR of ~ 100) to determine our

¹⁷The results for the posteriors obtained from DYNESTY [28,30] can be highly sensitive to the sampling settings. In this work, we follow other analyses in the literature [34] taking `n_live` = 2000 points as well as `n_act` = 10, and perform analytic phase marginalization which is crucial for stable results. We do not analytically marginalize over luminosity distance or merger time, although this should not affect the results. An additional technical point that is relevant here is that for waveform models with higher modes, it is known that this phase premarginalization step may break down if the $\ell > 2$ overtones give a significant contribution to the signal [81]. We have checked explicitly that running DYNESTY again without phase premarginalization (but with a larger `n_act` = 50) leads to identical posteriors on the well-measured parameters such as the chirp mass or luminosity distance, although with some instabilities in less well-constrained parameters such as the phase and the two tilts. Importantly, we note that PEREGRINE itself performs no such phase premarginalization and therefore the observed agreement is further indication that for the sources in question phase premarginalization is performing well here. Regardless, this highlights the fact that for comparisons, this check should be done on a case-by-case basis when working with higher-order mode waveform approximants.

¹⁸In general, this interesting source highlights the need for development in waveform modeling, as well as an investigation surrounding the implications of noise artifacts on inference [85,86]. This is additional motivation to develop flexible, targeted inference methods so that the impact of individual components of the setup (waveform approximant, noise model etc.) on the final inference can be reliably tested. This is very challenging to achieve with fully amortized methods.

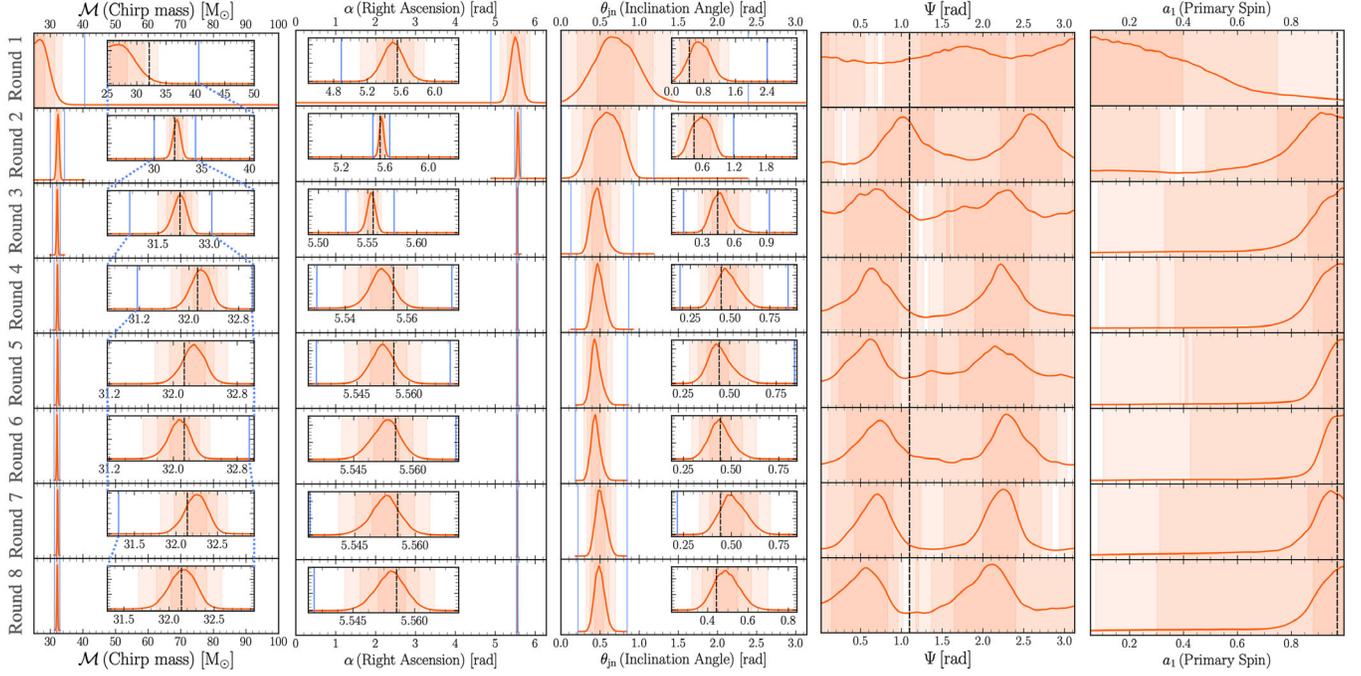


FIG. 3. Illustration of the truncation procedure as applied to the high SNR case study (C2). Each row corresponds to one round of inference for five separate parameters (chirp mass \mathcal{M} , right ascension α , inclination angle θ_{in} , polarization angle ψ , and primary spin a_1 , from left to right). The insets are zoomed in versions of the constrained region to illustrate the coverage of the true injected values (shown by black dashed lines). The solid blue lines show the prior truncation bounds derived during the inference step.

ability to achieve high precision inference. While this source is arguably slightly unphysical according to current merger rates (in terms of how loud it is), the aim of this test is to confirm that we can achieve extremely high precision inference through the application of the TMNRE algorithm.

The detector response for both of these target observations are shown in Fig. 2 above for the LIGO Hanford facility.

B. Parameter estimation with TMNRE

We perform parameter inference on both of the case studies using the implementation of the TMNRE algorithm in PEREGRINE. In particular, we fix the algorithm settings to the ones given in Table II. The key results of this section are shown in Figs. 4 and 5, while the truncation procedure is illustrated for a number of parameters in Fig. 3. A full set of 2d-posterior distributions can be found in the Appendix in Figs. 6 and 9.

Discussion. There are a number of levels to discuss these results. First, we consider the TMNRE algorithm and its applicability to gravitational wave data analysis. We can see clearly from the results in Figs. 4 and 5 that across all parameters we have excellent reconstruction of the injected parameters, including, e.g., the high primary spin which can be typically challenging to analyze [85,86]. More importantly, we see that we achieve very close agreement with the traditional nested sampling method DYNESTY [30], both in terms of accuracy of

parameter reconstruction, as well as the precision of the posteriors. To be quantitative, we compute the Jensen-Shannon divergence (JSD) between the two sets of posterior distributions. In Appendix B, we provide a full set of these values, averaged over all physical parameters across a number of different samplers including DYNESTY [30], but also CPNEST [87] and PTMCEE [26,88], as well as a complete description of our methodology for computing the JSD.¹⁹ Given these results, there are two key points to highlight here. First, we see that the variation between samplers as quantified by the JSD is an observation specific quantity and that in this case the variance is significantly larger for the high SNR case study (on average around $(10 - 30) \times 10^{-3}$ nat) compared to the low SNR one (on average around $(3 - 9) \times 10^{-3}$ nat). This trend follows at the level of variation between posterior distributions obtained from the same sampler (e.g., DYNESTY) run twice with different random seeds. Secondly, and more relevantly in terms of the PEREGRINE posteriors, we see that for both case studies, our results are fully consistent with the variations between individual samplers (although of course we will happen to show better agreement with one compared to another). In addition, it is worth noting that we also obtain significantly better coverage statistics than, e.g.,

¹⁹We use `n_live = 2000`, `n_act = 10` for `dynesty`, `n_live = 2000` for `cpnest` and `nsamples = 2000` for `ptmcee`.

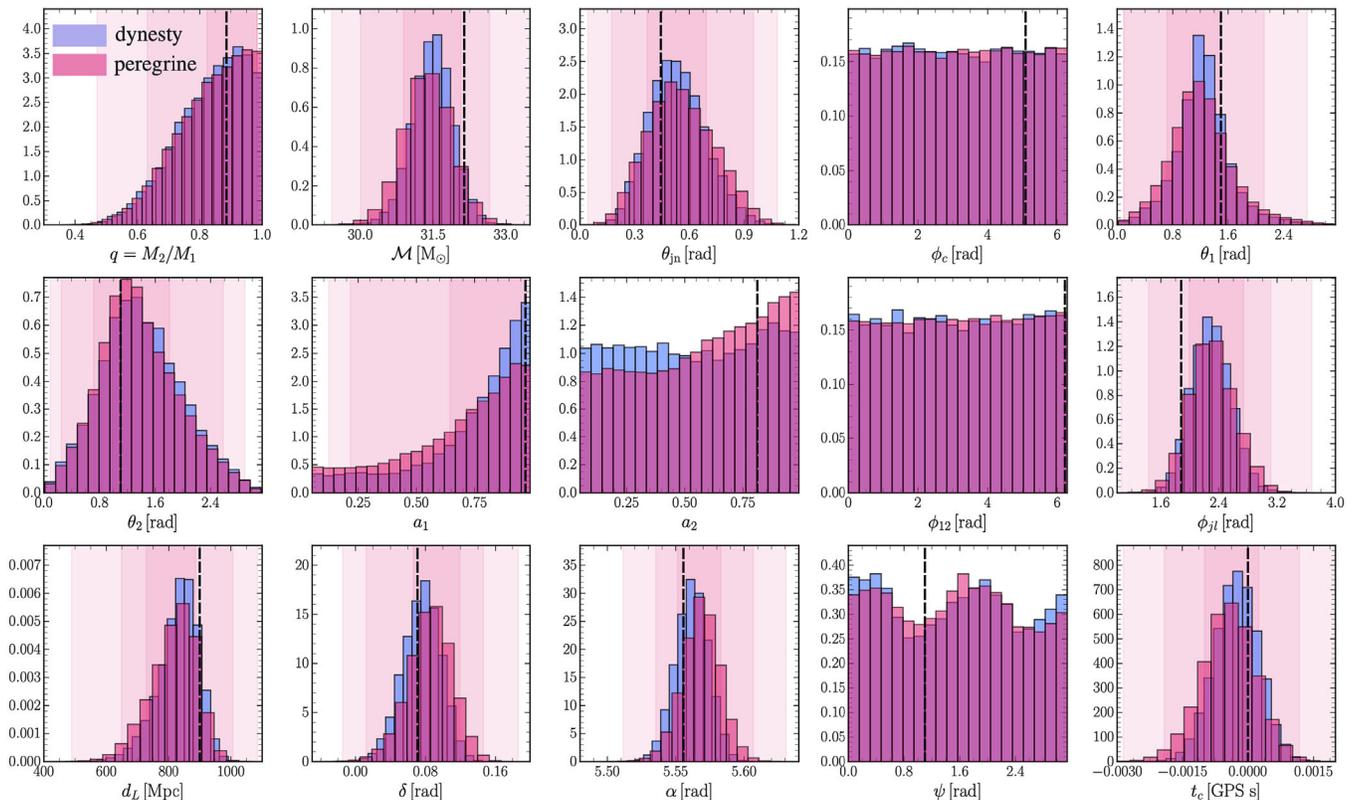


FIG. 4. Full set of 1d marginal posteriors for all parameters in the low SNR case study (C1). The inference results from the TMNRE algorithm are shown in pink, with the 1σ , 2σ , and 3σ contours shown underlaid behind. The purple histograms are the DYNESTY results for the analysis performed on the same observation (see Fig. 2). The black dashed lines indicate the true injected values.

BILBY-MCMC and DYNESTY (see Refs. [33,35] and the discussion below). As a final point, we can see that the main discrepancies are on parameters that are nonetheless poorly constrained, and certainly not measured to any precision—for example the secondary dimensionless spin a_2 , or the phase ϕ_c which is premarginalized in the sampling approach anyway.

A crucial element to achieving this level of agreement was to ensure that the same underlying noise realization, and corresponding observation was analyzed in both the TMNRE and nested sampling analysis. Performing the parameter inference on the same signal with a different noise realization will (correctly) lead to shifts in the posterior estimates of the parameter that only *on average* will be centred on the true value. This is something we actually take advantage of in the next section to perform additional quality checks of our inference. These are known broadly as coverage tests and are becoming common practice in the simulation-based inference literature to validate inference results [89,90].

Two of the important features of the TMNRE algorithm that we highlighted in Sec. II were the truncation and marginal elements. In Fig. 3, we highlight the former of these and illustrate the impact of truncation on various classes of parameters. For example, the first three

columns of panels illustrate the convergence of the posterior estimates over the rounds of truncation and re-simulation for parameters that are (extremely) well measured, such as the chirp mass or inclination angle. In the final two panels, we see another class of parameters (here ψ and a_1 as examples) which are not well constrained but are nonetheless reconstructed. We see for this second class of parameter that in the initial rounds, the marginal posterior estimate of, e.g., $p(\psi|x_0)$ is rather poor—and indeed, one should not interpret this as the true posterior, because the TMNRE algorithm has not converged. Once the other well measured parameters are tightly constrained, however, the quality of the inference of ψ and a_1 increases significantly, eventually converging to agree with the sampling result. This is a general trend when running TMNRE algorithms that convergence in other parameters over the course of the rounds ultimately leads to later improvements in posteriors for different parameters, see, e.g., [60]. It is also worth briefly commenting on multi-modality here: while the case studies we have chosen have well-constrained parameters such as the chirp mass that are single modes, multi-modal inference in gravitational waves is a common occurrence. In the context of PEREGRINE, the same level of efficiency could be realized in these scenarios by modifying how

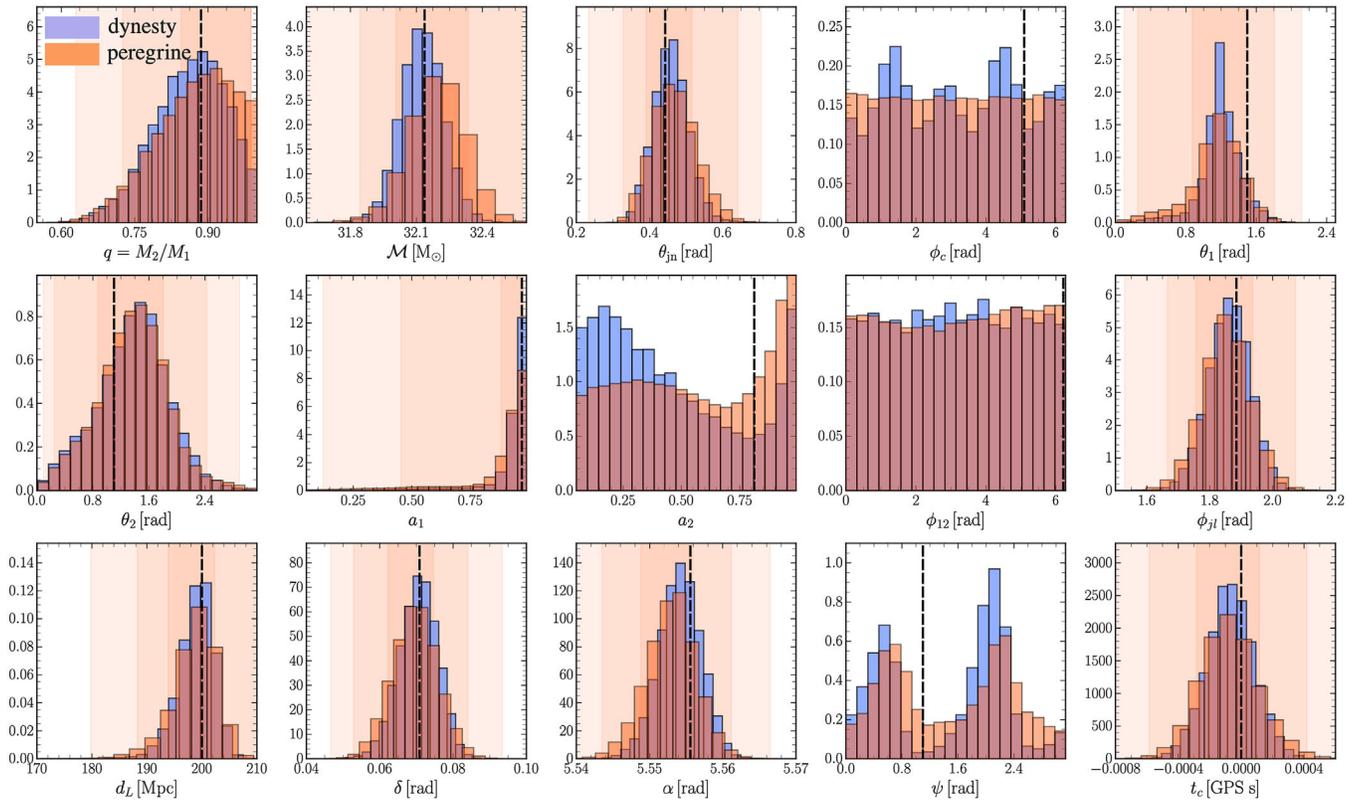


FIG. 5. Full set of 1d marginal posteriors for all parameters in the high SNR case study (C2). The inference results from the TMNRE algorithm are shown in orange, with the 1σ , 2σ , and 3σ contours shown underlaid behind. The purple histograms are the DYNesty results for the analysis performed on the same observation (see Fig. 2). The black dashed lines indicate the true injected values.

we perform the truncation of the prior proposal. Specifically, instead of bounding in some rectangular region, we could instead impose a threshold on the ratio and only sample in (potentially disconnected) regions of the prior that exceed this. This is currently under active development through, e.g., nested sampling applied to the prior distributions, and we emphasize it is just technical hurdle, rather than any limitation of the method. The second feature we highlighted was the benefit of directly estimating marginal posteriors $p(\vartheta|x_0)$ as opposed to first generating the full (N -dimensional) joint distribution. This plays a crucial role in simulation efficiency at no cost to statistical accuracy—indeed the comparison with nested sampling highlights that we obtain the correct marginal distributions even though we do not marginalize from a joint distribution.²⁰

We can explicitly highlight this simulation efficiency by comparing the number of waveform generation steps that were required in the nested sampling case versus in the TMNRE algorithm. In the low SNR case study, we

required around 720k waveform evaluations, while the corresponding DYNesty run required over 44 million. This represents an over 98% reduction in simulation budget for identical marginal posterior reconstruction. The difference is actually larger at the level of simulation time, since in each round of TMNRE, the waveform calls can be fully parallelized, while evaluations in nested sampling or MCMC broadly need to be performed sequentially. This parallelization is implemented as default in our PEREGRINE code. This is not an entirely fair comparison, since we are also required to train the classifiers at each stage. This should be compared to, e.g., the total likelihood evaluation time for a meaningful runtime comparison. Again, taking the low SNR case study as an example, the DYNesty run took around 43 hours versus around 12 hours for the TMNRE algorithm on the same architecture (including network training).²¹ For scenarios where the cost per simulation is much higher, however, this trade-off and benefit from simulation efficiency will become even more pronounced.

²⁰Of course, as discussed in Sec. II, we may not explicitly marginalize, but we do so implicitly by varying every parameter across all simulations.

²¹In future code versions, there is room for this time to be substantially improved through, e.g., accelerated waveform evaluations, reduction in network size, or a more efficient scheme for sampling from the truncated prior proposal distributions.

As a final note, while at the level of parameter estimation the 1d marginal posteriors are an informative set of results, it is often interesting to explore the parameter degeneracies. To do so, we take advantage of the flexibility of TMNRE to estimate only the marginal posteriors we are interested in testing and construct the full set of 2d posteriors. We make a comparison to the nested sampling contours in both case studies and find similarly good agreement. To achieve this, we required an additional 500k simulations in both the low and high SNR cases on which to train the set of marginal posterior classifiers. These are shown in the Appendix in Figs. 6 and 9.

C. Additional consistency checks

The results and comparisons to traditional methods discussed above are a helpful and meaningful way to test the data analysis pipeline that we have developed. On the other hand, the key motivation for developing PEREGRINE is to apply it to scenarios where the application of traditional methods is either impossible or extremely costly. This is slightly different motivation to current research efforts to perform, e.g., online parameter inference for LVKC data streams [64,91]. Given this, it is important to have additional ways to cross check and validate our inference results that do not rely on comparisons to sampling. This is an active area of ongoing research in simulation-based inference [48,89,92], however, a number of established techniques now exist.

The most common are known as coverage tests, which implement the idea discussed above that in the context of Bayesian parameter estimation, repeated inference over different noise/statistical realizations of the same signal should result in posteriors that shift relative to the true value. The rationale behind this sort of expected coverage test is that—simply as a result of statistical fluctuations—the $x\%$ credible interval for the posterior should contain the simulation-truth value $x\%$ of the time. For a more general discussion of this in the context of simulation-based inference, see, e.g., [89]. In practice, we can take a large sample²² of observations generated from the truncated prior (where we have converged posterior estimates), and perform inference on each observation. In each case we can note how often the injected value was contained inside the $x\%$ confidence interval and construct a cumulative distribution. A well-calibrated posterior distribution will be a totally diagonal line when the expected coverage is plotted against the empirical findings. Due to the fact that full inference must be done for each observation, this type of test is typically infeasible to do for sequential sampling methods such as MCMC. It should be noted, however, that this test is diagnostic in the sense that a failure implies a

poorly calibrated posterior, but success is not a definitive guarantee.

The coverage test results for all 15 parameters in both the high and low SNR case studies are provided in the Appendix (see Figs. 7, 8, 10 and 11). We see that across every parameter, we achieve extremely good posterior coverage, something that adds additional validation to our agreement with DYNESTY.

IV. CONCLUSIONS AND OUTLOOK

In this final section, we present the key conclusions of this work, before giving some outlook as to the possible use cases for our data analysis strategy. Ultimately the motivation for our method development comes from the analysis challenges facing the field of gravitational waves as we move toward the next generation of detectors [93–95]. Specifically, we will need methods that are highly scalable, simulation efficient, and flexible to take full advantage of future data. In this context, the key research contributions from this work are

- (i) *Simulation-based inference pipeline.* The key development put forward in this work is a sequential simulation-based inference pipeline for analyzing gravitational waves (see Sec. II for a discussion). Specifically, we implemented an algorithm known as truncated marginal neural ratio estimation (TMNRE) [50]. The motivation and advantages associated to this choice for gravitational waves are broadly centred around simulation efficiency and its applicability to individual observations. We showed for example that in a standard analysis for a binary black hole merger, we reduced the simulation budget (number of waveform evaluations) by over 98% compared to likelihood-based methods. This scaling is achieved primarily through the ability of the TMNRE algorithm to directly estimate marginal posteriors, and is the key argument for the use of our method in currently intractable sampling problems such as the analysis of overlapping sources [38–42].
- (ii) *Case studies.* We studied two cases where we analyzed highly spinning precessing binary black hole mergers, one with a high SNR of around 100, and another with a relatively low SNR of approximately 20. In each case, we performed full parameter inference using our method on observations generated using the state-of-the-art SEOBNRv4 PHM waveform approximant [21] (see Fig. 2), although the method can be applied to any waveform model. We compared our results for the 1d- and 2d-marginal posteriors to posterior samples obtained from nested sampling and demonstrated the excellent agreement across all parameters (see Figs. 4 and 5). In addition, we carried out expected coverage tests to validate the behavior

²²In the tests shown below, we generated 2000 extra simulations from the final round truncated prior region.

of our posterior estimates (see, e.g., Figs. 10 and 11 for the high SNR case).

- (iii) *Robust and flexible method.* Our implementation of TMNRE is not the first method to apply modern machine learning techniques to the analysis of gravitational waves [59,64,71,96–99]. On the other hand, these other impressive parameter estimation pipelines typically have a different goal in mind—rapid or instantaneous parameter inference. This has great relevance when attempting to, e.g., optimize signal triggering or follow up quickly on electromagnetic counterparts. However, this comes at the cost of complicated architectures, or bespoke likelihood designs [91,100,101] which render these approaches less flexible for studying arbitrary classes of signals. In contrast, our motivation is to develop a method that is robust and flexible enough to be used to analyze a wide range of signal classes in specific observations. In order to study a new class of signal—only a new forward simulation model is required, while the rest of the pipeline will remain identical. Importantly, in cases where traditional methods are not available, the coverage tests described in Sec. III will still be applicable, so we can still validate our results.
- (iv) *Code release.* Alongside this work, we have developed an implementation of our analysis method, known as PEREGRINE (built on top of SWYFT [49,50]). This is a highly modular and scalable library that allows the user to implement their own `Simulator` class. Given the simulation efficiency and relatively small network size, it is usable on an individual basis with reasonable access to consumer hardware. Both the algorithm as well as the coverage tests are fully developed and tested, with examples available. Access to PEREGRINE can be requested at this Ref. [102].

A. Outlook

As discussed in the introduction to this work, we are currently in an exciting era of gravitational wave experimental development. With future detectors such as the Einstein Telescope [93] or LISA [95] taking shape, the future gravitational wave sky could be extremely loud, varied, and informative. On the other hand, there are clear data analysis challenges associated to fully utilizing this data [67,68]. Two important examples of these challenges include (i) the analysis of two (or more) overlapping gravitational wave signals, which will arise as a result of large increases in detector sensitivity [38–42,103], and (ii) the identification and analysis of stochastic signals

that will be relevant to space-based observatories such as LISA [95]. In the former case, the challenge is a sampling one, with current methods potentially taking months to analyze a single event. In the latter case, there is also potentially a large sampling challenge in an attempt to fit multiple signal templates or noise models, but there are also statistical challenges that could arise from correlated signal or detector noise, as well as significant simulation costs.

It is this class of data analysis challenges that we had in mind when developing PEREGRINE.²³ There are a number of reasons to believe that the framework put forward here can approach these problems and extract the maximum science/physics results. Firstly, the method is designed to be highly simulation efficient, something we have demonstrated explicitly in the case of binary black hole mergers where we saw a $\sim 98\%$ reduction in simulation budget. This is one of the keys to solving sampling based problems, where we can directly perform inference at the level of marginal distributions rather than solve the full joint problem. Assuming the same simulation efficiency holds, this could lead to a reduction in analysis time for a problem such as parameter inference for multiple overlapping waveforms from $\mathcal{O}(\text{months})$ [41] to less than a day. Similarly, TMNRE and other simulation-based inference approaches are by construction implicit-likelihood methods. This is equivalent to the statement that one only needs a forward simulator to perform inference, rather than an explicit likelihood. In the case of, e.g., possible correlated noise in the stochastic gravitational wave background [43,44], this feature could unlock a number of physics questions that are currently intractable. It is also worth noting that nonstationary noise, which is typically challenging to deal with in a likelihood-based framework is simple in the context of simulation-based inference and PEREGRINE. In particular, the analysis would remain identical if, e.g., the PSD was time-dependent, provided this was included in the detector aspect of the simulator. We will leave these to future works to analyze these situations in detail.

To summarize, it is clear that in the era of high-sensitivity gravitational wave experiments, flexible and scalable data analysis strategies will be crucial to uncovering exciting fundamental physics discoveries. In this

²³Although in addition, motivated by the observed simulation efficiency, it would be interesting to explore the extent to which the inference can be fully amortized by significantly expanding the initial simulation budget, after which approximately real time inference could be performed. We will leave this aspect to future work, however, along with investigations as to how to include the effect of nonstationary noise in such an approach.

work, we have presented PEREGRINE as a community-based tool that can hopefully take some steps toward that goal.

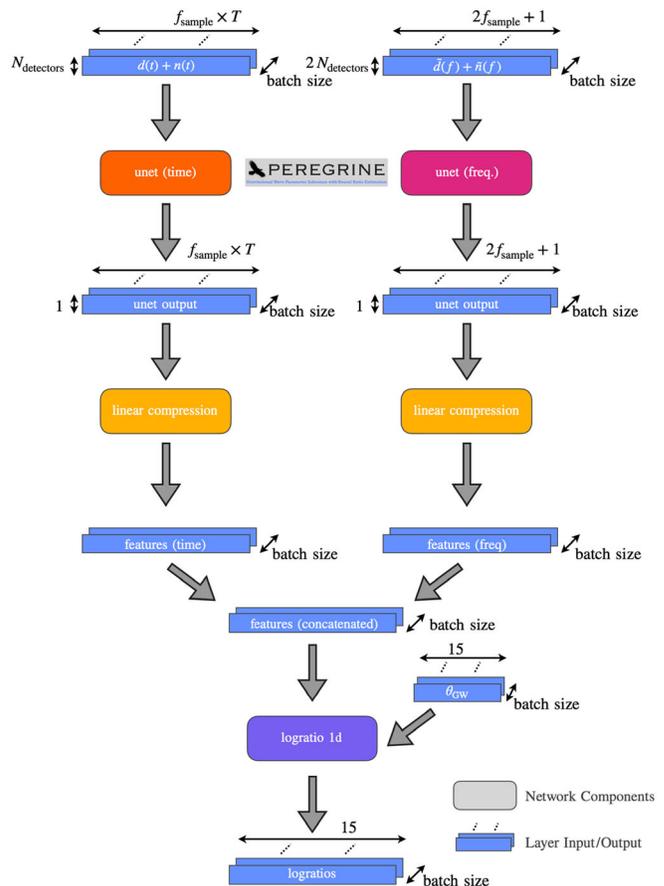
Code and data availability. See GitHub in Ref. [104] for the PEREGRINE analysis and inference library.

ACKNOWLEDGMENTS

This work is part of the project CORTEX (NWA .1160.18.316) of the research programme NWA-ORC which is (partly) financed by the Dutch Research Council (NWO). Additionally, C. W. acknowledges funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 864035). U. B. is supported through the CORTEX project of the NWA-ORC with project number NWA.1160.18.316 which is partly financed by the Dutch Research Council (NWO). J. A. is supported through the research program “The Hidden Universe of Weakly Interacting Particles” with project number 680.92.18.03 (NWO Vrije Programma), which is partly financed by the Nederlandse Organisatie voor Wetenschappelijk Onderzoek (Dutch Research Council). S. N. acknowledges support from the NWO Projectruimte grant (Samaya Nissanke). B. K. M. is funded by the University of Amsterdam Faculty of Science (FNWI), Informatics Institute (IvI), and the Institute of Physics (IoP). B. K. M. is also affiliated with the European Laboratory for Learning and Intelligent Systems (ELLIS Society). The main analysis for this work was carried out on the Lisa and Snellius Computing Clusters at SURFsara. We are very grateful to Thomas Edwards and Sam Witte for their careful reading of the draft and useful comments. We also thank Andrew Williamson and Pablo Bosch for helpful conversations about this project.

APPENDIX A: NETWORK ARCHITECTURE

The network architecture used for the analysis in PEREGRINE is shown in the figure below. As described in the text in Sec. II, we use the flexibility of the TMNRE implementation to automatically extract a small set of summary statistics from both the time and frequency domain strains after processing and compressing the input data. This summary can then be passed to the ratio estimator along with the physical parameters θ_{GW} to perform the inference. The full PyTorch [105] implementation of the network components shown below (specifically the unet and linear compression networks) can be found in the `InferenceNetwork` class of PEREGRINE.



APPENDIX B: COMPARISON WITH LIKELIHOOD BASED SAMPLERS

In order to validate our results in Sec. III, we quantitatively compare the PEREGRINE posteriors against those obtained using traditional sampling methods such as nested sampling

TABLE III. Average Jensen-Shannon divergences over all physical parameters $\theta_{\text{GW}} = (q, \mathcal{M}, \dots)$ for different sampler combinations. Values below the diagonal (in pink boxes) refer to the low SNR case study (C1), while those above the diagonal (in orange boxes) are for the high SNR case (C2).

JSD [10^{-3} nat]	PEREGRINE	DYNESTY	DYNESTY (re-run)	PTEMCEE	CPNEST
PEREGRINE		18.0	13.1	21.0	20.6
DYNESTY	9.75		4.02	10.0	35.2
DYNESTY (re-run)	10.9	0.55		10.6	29.9
PTEMCEE	16.5	7.01	7.56		35.9
CPNEST	6.82	3.52	4.46	8.60	

and MCMC. Specifically, we compute the Jensen-Shannon divergence among the 1d posteriors obtained using PEREGRINE, DYNESTY, PTEMCEE, and CPNEST. Our implementation including all the data relevant to this article are

available publicly on our Zenodo page [106], along with files, along with files that can regenerate the JSD values (both on average and on a parameter-by-parameter basis) given below in Table III.

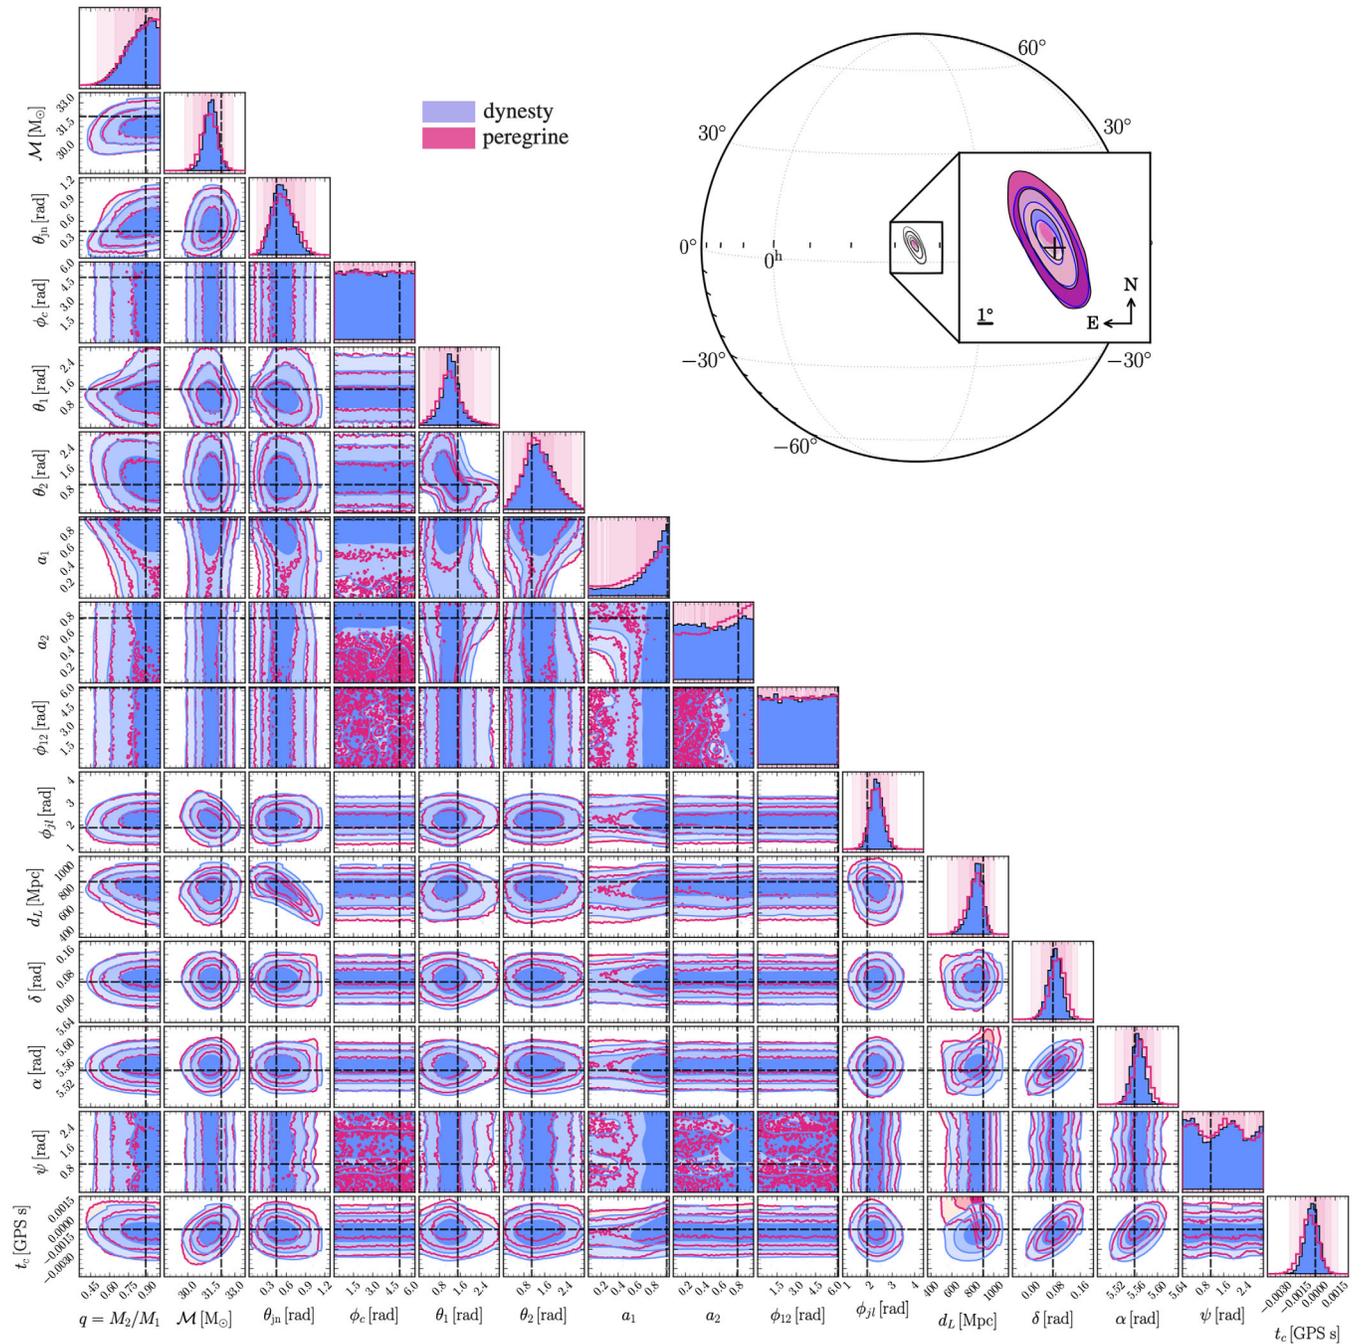


FIG. 6. Full set of 2d marginal posteriors in the low SNR case study (C1). The results from the TMNRE analysis are shown in pink, while the corresponding DYNESTY analysis is shown in blue. The darker and subsequently lighter contours in the 2d marginals indicate the 1σ , 2σ and 3σ confidence intervals respectively. The sky map shows the (α, δ) contours centered at the injection value.

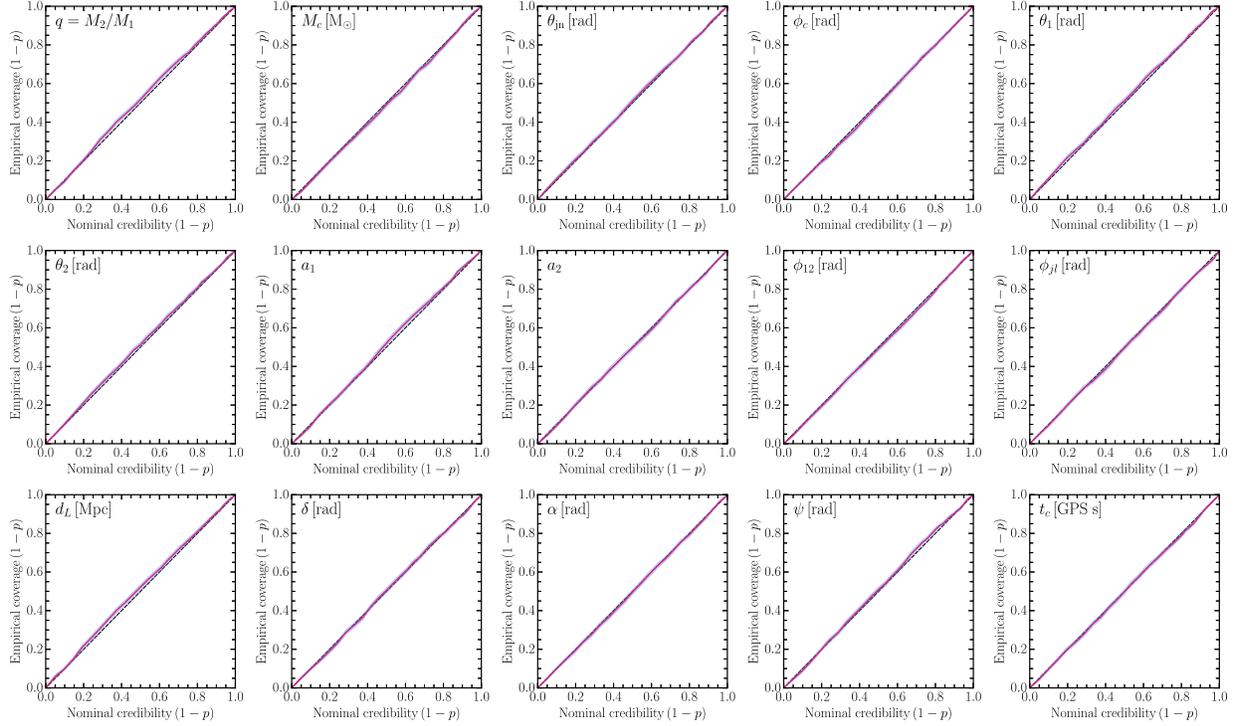


FIG. 7. Coverage results for the low SNR case study (C1) for all parameters. This compares the expected coverage of the true value as a percentage on the x -axis against the actual coverage of our ratio estimator on the y -axis. The pink line indicates the average coverage, while the blue contour indicates the 68% confidence interval on the coverage.

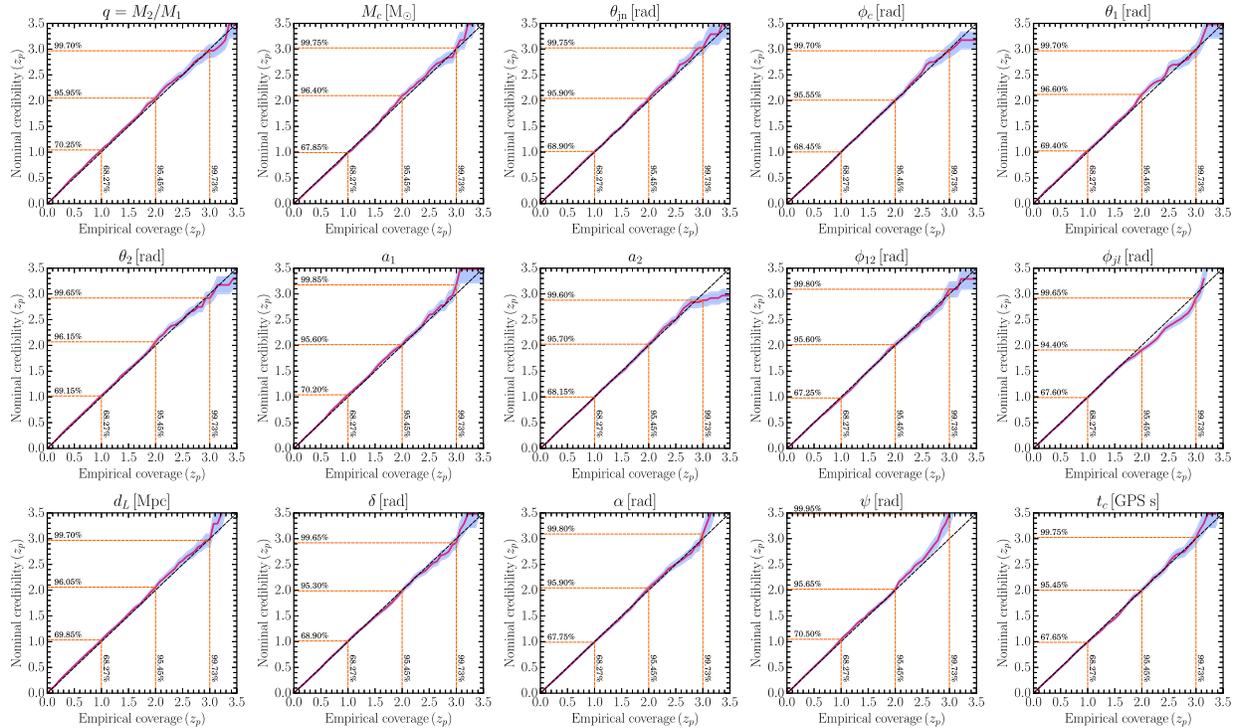


FIG. 8. Coverage results for the low SNR case study (C1) for all parameters. This is the same information as Fig. 10, but with z_p defined by $p = \int_{-z_p}^{z_p} dz / \sqrt{2\pi} \exp(-z^2/2)$. This places more emphasis on the behavior of the posteriors in the tail regions. The pink line indicates the average coverage, while the blue contour indicates the 68% confidence interval on the coverage.

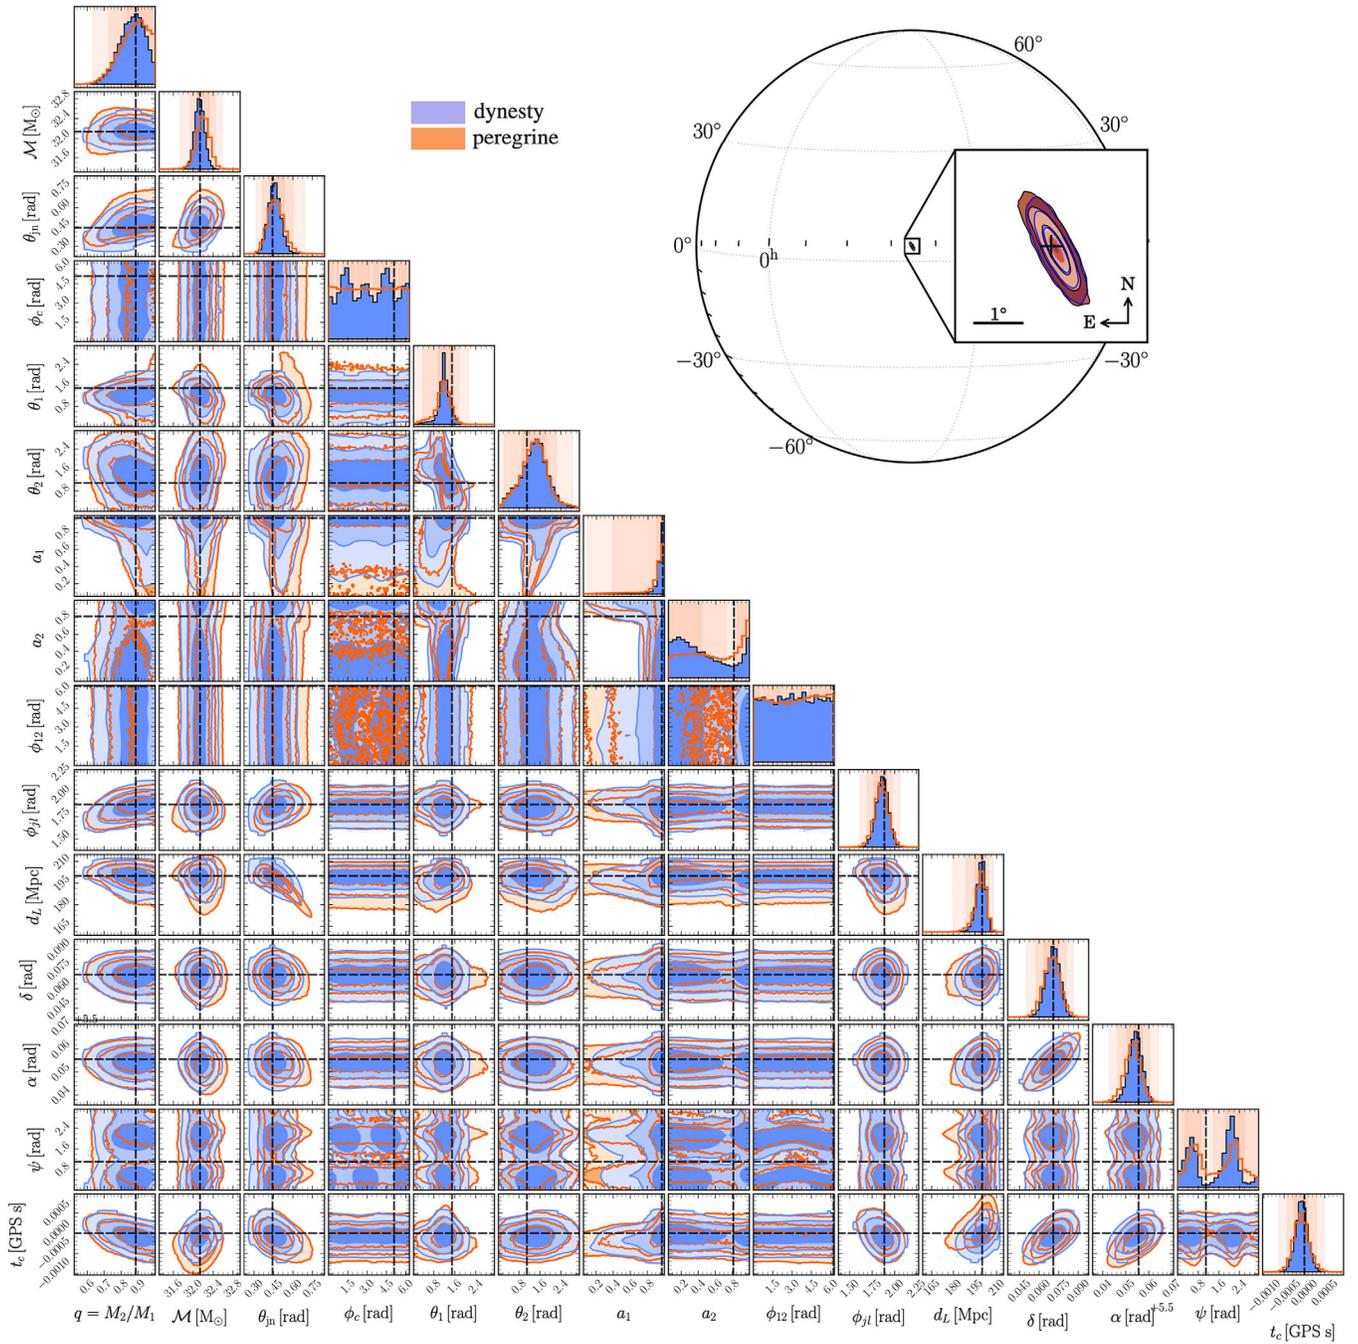


FIG. 9. Full set of 2d marginal posteriors in the high SNR case study (C2). The results from the TMNRE analysis are shown in orange, while the corresponding DYNESTY analysis is shown in blue. The darker and subsequently lighter contours in the 2d marginals indicate the 1σ , 2σ , and 3σ confidence intervals respectively. The sky map shows the (α, δ) contours centered at the injection value.

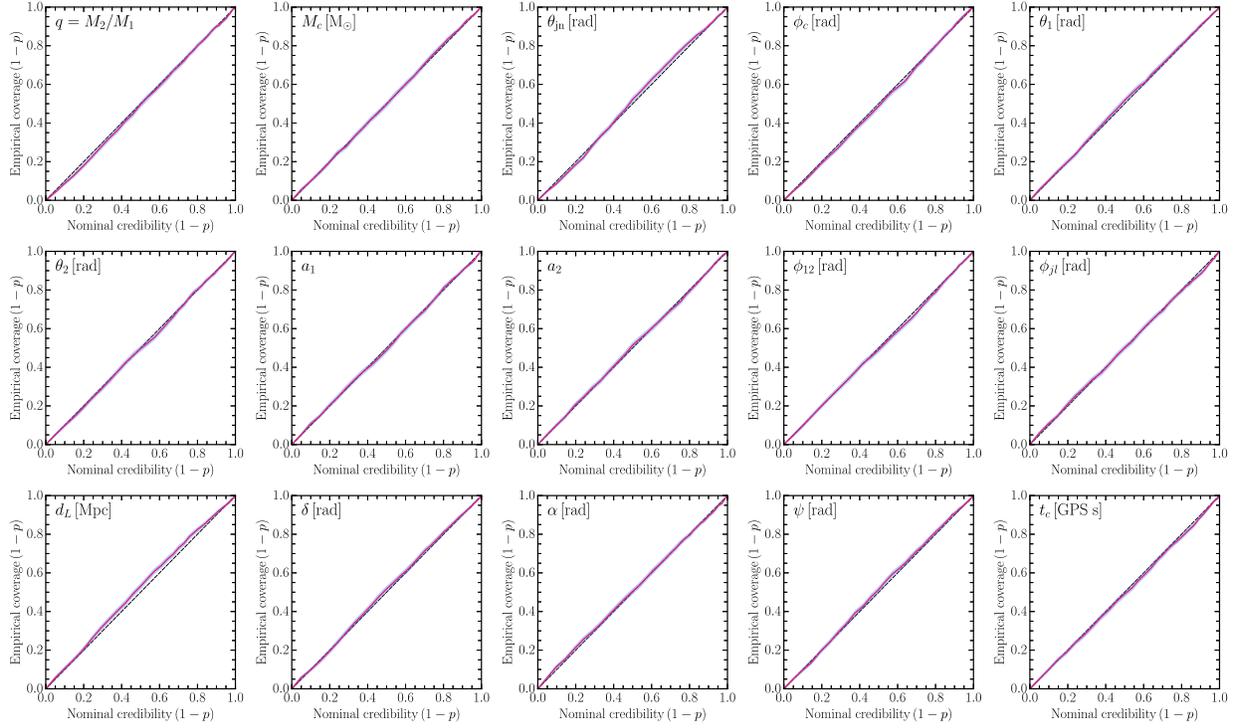


FIG. 10. Coverage results for the high SNR case study (C2) for all parameters. This compares the expected coverage of the true value as a percentage on the x -axis against the actual coverage of our ratio estimator on the y -axis. The pink line indicates the average coverage, while the blue contour indicates the 68% confidence interval on the coverage.

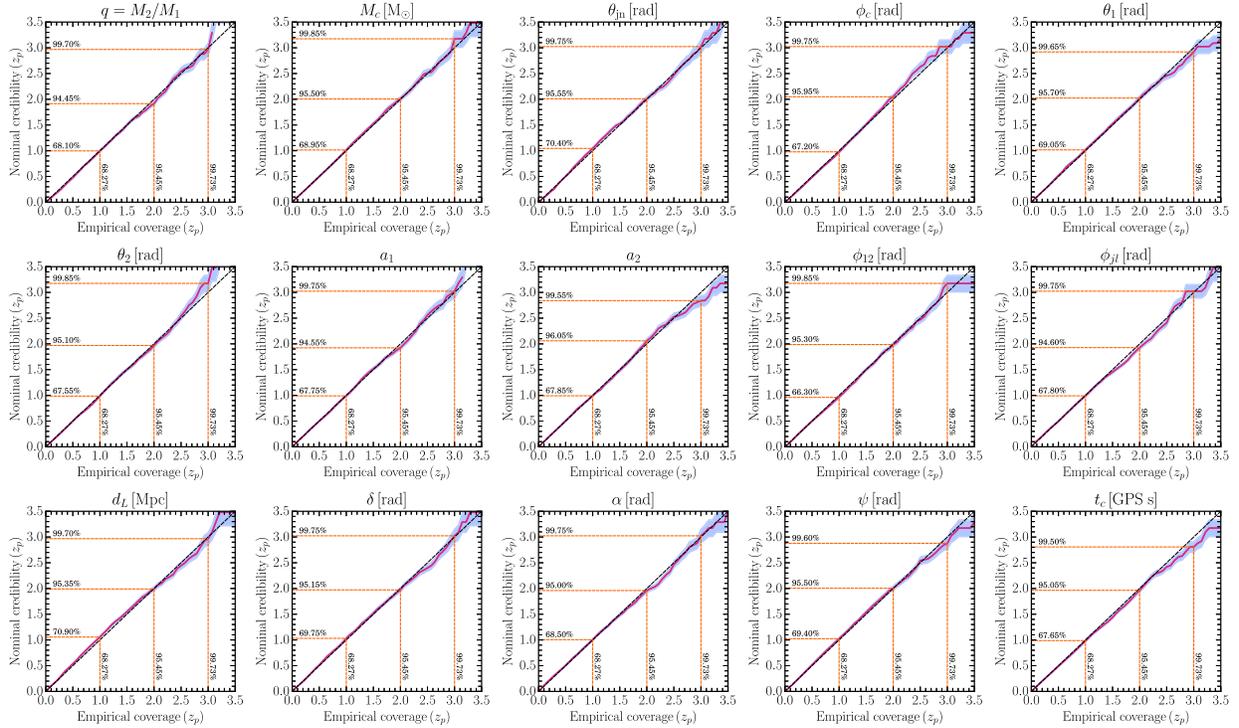


FIG. 11. Coverage results for the high SNR case study (C2) for all parameters. This is the same information as Fig. 10, but with z_p defined by $p = \int_{-z_p}^{z_p} dz \, 1/\sqrt{2\pi} \exp(-z^2/2)$. This places more emphasis on the behavior of the posteriors in the tail regions. The pink line indicates the average coverage, while the blue contour indicates the 68% confidence interval on the coverage.

- [1] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **116**, 061102 (2016).
- [2] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. D* **103**, 122002 (2021).
- [3] B. P. Abbott *et al.* (LIGO Scientific, Virgo, and VIRGO Collaborations), *Astrophys. J.* **909**, 218 (2021).
- [4] S. Nissanke, D. E. Holz, S. A. Hughes, N. Dalal, and J. L. Sievers, *Astrophys. J.* **725**, 496 (2010).
- [5] S. Nissanke, J. Sievers, N. Dalal, and D. Holz, *Astrophys. J.* **739**, 99 (2011).
- [6] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Astrophys. J. Lett.* **913**, L7 (2021).
- [7] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. Lett.* **121**, 161101 (2018).
- [8] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), [arXiv:2108.01045](https://arxiv.org/abs/2108.01045).
- [9] A. H. Nitz, C. D. Capano, S. Kumar, Y.-F. Wang, S. Kastha, M. Schäfer, R. Dhurkunde, and M. Cabero, *Astrophys. J.* **922**, 76 (2021).
- [10] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, *Phys. Rev. D* **100**, 023011 (2019).
- [11] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, *Phys. Rev. D* **101**, 083030 (2020).
- [12] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. X* **11**, 021053 (2021).
- [13] R. Abbott *et al.* (LIGO Scientific, VIRGO, and KAGRA Collaborations), [arXiv:2111.03606](https://arxiv.org/abs/2111.03606) [*Phys. Rev. X* (to be published)].
- [14] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, *Phys. Rev. D* **85**, 122006 (2012).
- [15] M. Isi, R. Smith, S. Vitale, T. J. Massinger, J. Kanner, and A. Vajpeyi, *Phys. Rev. D* **98**, 042007 (2018).
- [16] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, *Phys. Rev. D* **100**, 023011 (2019).
- [17] B. P. Abbott *et al.* (KAGRA, LIGO Scientific, Virgo, and VIRGO Collaborations), *Living Rev. Relativity* **21**, 3 (2018).
- [18] P. Petrov, L. P. Singer, M. W. Coughlin, V. Kumar, M. Almualla, S. Anand, M. Bulla, T. Dietrich, F. Foucart, and N. Guessoum, *Astrophys. J.* **924**, 54 (2022).
- [19] C. Cutler and E. E. Flanagan, *Phys. Rev. D* **49**, 2658 (1994).
- [20] E. Thrane and C. Talbot, *Publ. Astron. Soc. Pac.* **36**, e010 (2019).
- [21] S. Ossokine *et al.*, *Phys. Rev. D* **102**, 044055 (2020).
- [22] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. Jiménez Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044007 (2016).
- [23] G. Pratten, S. Husa, C. Garcia-Quiros, M. Colleoni, A. Ramos-Buades, H. Estelles, and R. Jaume, *Phys. Rev. D* **102**, 064001 (2020).
- [24] T. D. P. Edwards, K. W. K. Wong, K. K. H. Lam, A. Coogan, D. Foreman-Mackey, M. Isi, and A. Zimmerman, [arXiv:2302.05329](https://arxiv.org/abs/2302.05329).
- [25] D. Mackay, *Informaiton Theory, Inference, and Learning Algorithms* (Cambridge University Press, Cambridge, England, 2003).
- [26] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
- [27] J. Skilling, *Bayesian Anal.* **1**, 833 (2006).
- [28] W. J. Handley, M. P. Hobson, and A. N. Lasenby, *Mon. Not. R. Astron. Soc.* **453**, 4384 (2015).
- [29] G. Ashton *et al.*, *Nature (London)* **2** (2022), 39.
- [30] J. S. Speagle, *Mon. Not. R. Astron. Soc.* **493**, 3132 (2020).
- [31] J. Veitch *et al.*, *Phys. Rev. D* **91**, 042003 (2015).
- [32] C. M. Biwer, C. D. Capano, S. De, M. Cabero, D. A. Brown, A. H. Nitz, and V. Raymond, *Publ. Astron. Soc. Pac.* **131**, 024503 (2019).
- [33] G. Ashton *et al.*, *Astrophys. J. Suppl. Ser.* **241**, 27 (2019).
- [34] I. M. Romero-Shaw *et al.*, *Mon. Not. R. Astron. Soc.* **499**, 3295 (2020).
- [35] G. Ashton and C. Talbot, *Mon. Not. R. Astron. Soc.* **507**, 2037 (2021).
- [36] M. J. Williams, J. Veitch, and C. Messenger, *Mach. Learn. Sci. Tech.* **4**, 035011 (2023).
- [37] F. Iacovelli, M. Mancarella, S. Foffa, and M. Maggiore, *Astrophys. J.* **941**, 208 (2022).
- [38] E. Pizzati, S. Sachdev, A. Gupta, and B. Sathyaprakash, *Phys. Rev. D* **105**, 104016 (2022).
- [39] A. Antonelli, O. Burke, and J. R. Gair, *Mon. Not. R. Astron. Soc.* **507**, 5069 (2021).
- [40] A. Samajdar, J. Janquart, C. Van Den Broeck, and T. Dietrich, *Phys. Rev. D* **104**, 044003 (2021).
- [41] J. Janquart, T. Baka, A. Samajdar, T. Dietrich, and C. Van Den Broeck, [arXiv:2211.01304](https://arxiv.org/abs/2211.01304).
- [42] J. Langendorff, A. Kolmus, J. Janquart, and C. Van Den Broeck, *Phys. Rev. Lett.* **130**, 171402 (2023).
- [43] C. Caprini and D. G. Figueroa, *Classical Quantum Gravity* **35**, 163001 (2018).
- [44] N. Christensen, *Rep. Prog. Phys.* **82**, 016903 (2019).
- [45] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. D* **100**, 061101 (2019).
- [46] K. Cranmer, J. Brehmer, and G. Louppe, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 30055 (2020).
- [47] J. Brehmer and K. Cranmer, [arXiv:2010.06439](https://arxiv.org/abs/2010.06439).
- [48] J.-M. Lueckmann, J. Boelts, D. S. Greenberg, P. J. Gonçalves, and J. H. Macke, [arXiv:2101.04653](https://arxiv.org/abs/2101.04653).
- [49] B. K. Miller, A. Cole, P. Forré, G. Louppe, and C. Weniger, in *Proceedings of the 35th Conference on Neural Information Processing Systems* (2021), [arXiv:2107.01214](https://arxiv.org/abs/2107.01214).
- [50] B. K. Miller, A. Cole, C. Weniger, F. Nattino, O. Ku, and M. W. Grootes, *J. Open Source Software* **7**, 4205 (2022).
- [51] A. Tejero-Cantero, J. Boelts, M. Deistler, J.-M. Lueckmann, C. Durkan, P. J. Gonçalves, D. S. Greenberg, and J. H. Macke, *J. Open Source Software* **5**, 2505 (2020).
- [52] J. Alsing, T. Charnock, S. Feeney, and B. Wandelt, *Mon. Not. R. Astron. Soc.* **488**, 4440 (2019).
- [53] A. Cole, B. K. Miller, S. J. Witte, M. X. Cai, M. W. Grootes, F. Nattino, and C. Weniger, *J. Cosmol. Astropart. Phys.* **09** (2022) 004.
- [54] N. A. Montel, A. Coogan, C. Correa, K. Karchev, and C. Weniger, *Mon. Not. R. Astron. Soc.* **518**, 2746 (2022).
- [55] N. Anau Montel and C. Weniger, in *Proceedings of the 36th Conference on Neural Information Processing Systems* (2022), [arXiv:2211.04291](https://arxiv.org/abs/2211.04291).
- [56] T. L. Mäkinen, T. Charnock, J. Alsing, and B. D. Wandelt, *J. Cosmol. Astropart. Phys.* **11** (2021) 049.

- [57] A. Dimitriou, C. Weniger, and C. A. Correa, [arXiv:2206.11312](#).
- [58] S. Gagnon-Hartman, J. Ruan, and D. Haggard, *Mon. Not. R. Astron. Soc.* **520**, 1 (2023).
- [59] A. Delaunoy, A. Wehenkel, T. Hinderer, S. Nissanke, C. Weniger, A. R. Williamson, and G. Louppe, [arXiv:2010.12931](#).
- [60] K. Karchev, R. Trotta, and C. Weniger, *Mon. Not. R. Astron. Soc.* **520**, 1056 (2023).
- [61] K. Lin, M. von Wietersheim-Kramsta, B. Joachimi, and S. Feeney, [arXiv:2212.04521](#).
- [62] M. Dax, S. R. Green, J. Gair, M. Deistler, B. Schölkopf, and J. H. Macke, [arXiv:2111.13139](#).
- [63] J. Wildberger, M. Dax, S. R. Green, J. Gair, M. Pürrer, J. H. Macke, A. Buonanno, and B. Schölkopf, *Phys. Rev. D* **107**, 084046 (2023).
- [64] M. Dax, S. R. Green, J. Gair, J. H. Macke, A. Buonanno, and B. Schölkopf, *Phys. Rev. Lett.* **127**, 241103 (2021).
- [65] G. Pratten *et al.*, *Phys. Rev. D* **103**, 104056 (2021).
- [66] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, *Phys. Rev. Lett.* **113**, 151101 (2014).
- [67] S. Babak *et al.* (Mock LISA Data Challenge Task Force Collaboration), *Classical Quantum Gravity* **27**, 084009 (2010).
- [68] M. B. Schäfer *et al.*, *Phys. Rev. D* **107**, 023021 (2023).
- [69] G. Papamakarios and I. Murray, [arXiv:1605.06376](#).
- [70] J. Zeghal, F. Lanasse, A. Boucaud, B. Remy, and E. Aubourg, [arXiv:2207.05636](#).
- [71] M. Dax, S. R. Green, J. Gair, M. Pürrer, J. Wildberger, J. H. Macke, A. Buonanno, and B. Schölkopf, *Phys. Rev. Lett.* **130**, 171403 (2023).
- [72] F. Rozet and G. Louppe, [arXiv:2110.00449](#).
- [73] A. Delaunoy, J. Hermans, F. Rozet, A. Wehenkel, and G. Louppe, [arXiv:2208.13624](#).
- [74] B. K. Miller, C. Weniger, and P. Forré, [arXiv:2210.06170](#).
- [75] J. Hermans, V. Begy, and G. Louppe, [arXiv:1903.04057](#).
- [76] C. Durkan, I. Murray, and G. Papamakarios, [arXiv:2002.03712](#).
- [77] O. Ronneberger, P. Fischer, and T. Brox, [arXiv:1505.04597](#).
- [78] M. Sunnaker, A. Busetto, E. Numminen, J. Corander, M. Foll, and C. Dessimoz, *PLoS Comput. Biol.* **9**, e1002803 (2013).
- [79] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Classical Quantum Gravity* **37**, 055002 (2020).
- [80] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. X* **9**, 031040 (2019).
- [81] E. Thrane and C. Talbot, *Pub. Astron. Soc. Aust.* **36**, e010 (2019).
- [82] M. Hannam *et al.*, *Nature (London)* **610**, 652 (2022).
- [83] E. Payne, S. Hourihane, J. Golomb, R. Udall, R. Udall, D. Davis, and K. Chatziioannou, *Phys. Rev. D* **106**, 104017 (2022).
- [84] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. J. Forteza, and A. Bohé, *Phys. Rev. D* **93**, 044006 (2016).
- [85] I. M. Romero-Shaw, D. Gerosa, and N. Loutrel, *Mon. Not. R. Astron. Soc.* **519**, 5352 (2023).
- [86] Y. Xu and E. Hamilton, *Phys. Rev. D* **107**, 103049 (2023).
- [87] J. Veitch, W. D. Pozzo, A. Lyttle, M. Williams, C. Talbot, M. Pitkin, G. Ashton, Cody, M. Hübner, A. Nitz, D. Macleod, G. Carullo, G. Davies, and Tony, [cpnest: v0.11.3 \(2021\)](#), <https://zenodo.org/record/7385221>.
- [88] W. D. Voudsen, W. M. Farr, and I. Mandel, *Mon. Not. R. Astron. Soc.* **455**, 1919 (2016).
- [89] J. Hermans, A. Delaunoy, F. Rozet, A. Wehenkel, V. Begy, and G. Louppe, [arXiv:2110.06581](#).
- [90] P. Lemos, A. Coogan, Y. Hezaveh, and L. Perreault-Levasseur, [arXiv:2302.03026](#).
- [91] K. W. K. Wong, M. Isi, and T. D. P. Edwards, [arXiv:2302.05333](#).
- [92] P. Cannon, D. Ward, and S. M. Schmon, [arXiv:2209.01845](#).
- [93] M. Maggiore *et al.*, *J. Cosmol. Astropart. Phys.* **03** (2020) 050.
- [94] D. Reitze *et al.*, *Bull. Am. Astron. Soc.* **51**, 035 (2019).
- [95] P. Amaro-Seoane, H. Audley, S. Babak, J. Baker, E. Barausse, P. Bender, E. Berti *et al.*, [arXiv:1702.00786](#).
- [96] H. Gabbard, C. Messenger, I. S. Heng, F. Tonolini, and R. Murray-Smith, *Nat. Phys.* **18**, 112 (2022).
- [97] S. R. Green and J. Gair, *Mach. Learn. Sci. Tech.* **2**, 03LT01 (2021).
- [98] A. J. K. Chua and M. Vallisneri, *Phys. Rev. Lett.* **124**, 041102 (2020).
- [99] J. a. D. Álvares, J. A. Font, F. F. Freitas, O. G. Freitas, A. P. Morais, S. Nunes, A. Onofre, and A. Torres-Forné [arXiv:2011.10425](#).
- [100] B. Zackay, L. Dai, and T. Venumadhav, [arXiv:1806.08792](#).
- [101] N. Leslie, L. Dai, and G. Pratten, *Phys. Rev. D* **104**, 123030 (2021).
- [102] <https://github.com/PEREGRINE-GW/peregrine/>.
- [103] R. Busicchio, E. Roebber, J. M. Goldstein, and C. J. Moore, *Phys. Rev. D* **100**, 084041 (2019).
- [104] Please see GitHub at <https://github.com/PEREGRINE-GW/peregrine/> for data availability.
- [105] A. Paszke, S. Gross, F. Massa *et al.*, [arXiv:1912.01703](#).
- [106] <https://zenodo.org/record/7788596>.