

---

# Normalizing Flows for Hierarchical Bayesian Analysis: A Gravitational Wave Population Study

---

**David Ruhe**

AI4Science Lab, AMLab, Anton Pannekoek Institute  
University of Amsterdam  
d.ruhe@uva.nl

**Kaze Wong**

Center for Computational Astrophysics  
Flatiron Institute  
kwong@flatironinstitute.org


**Miles Cranmer**

Department of Astrophysical Sciences  
Princeton University  
mcranmer@princeton.edu

**Patrick Forré**

AI4Science Lab, AMLab  
University of Amsterdam  
p.d.forre@uva.nl

## Abstract

We propose parameterizing the population distribution of the gravitational wave population modeling framework (Hierarchical Bayesian Analysis) with a normalizing flow. We first demonstrate the merit of this method on illustrative experiments and then analyze four parameters of the latest LIGO data release: primary mass, secondary mass, redshift, and effective spin. Our results show that despite the small and notoriously noisy dataset, the posterior predictive distributions (assuming a prior over the parameters of the flow) of the observed gravitational wave population recover structure that agrees with robust previous phenomenological modeling results while being less susceptible to biases introduced by less-flexible distribution models. Therefore, the method forms a promising flexible, reliable replacement for population inference distributions, even when data is highly noisy. 

## 1 Introduction

The number of detected gravitational wave events is rapidly increasing [1, 31, 51, 32, 45, 2, 3]. Consequently, datasets of inferred properties such as mass, spin, and redshift have become feasible, the most recent being produced from GWTC-3 [3]. At this point, there is sufficient data for population-level analyses [26, 14, 50, 39, 40, 18, 21, 22, 38, 42, 5, 4]. Many previous methods assume a simplistic phenomenological shape of the distribution such as a power law. However, as indicated in [42], it is crucial that the proposed model class can accurately represent the data; otherwise, the reconstruction may miss essential population-level properties and correlations. An imposed phenomenological shape can suffer from human-induced bias, leading to potentially incorrect conclusions. In addition, the detection rate increases steeply with improved detector sensitivity, and it will be difficult to extract all the information from growing datasets. [42] move away from this approach by proposing a mixture model of weighted Gaussians (and a power law) that is expected to be capable of modeling a variety of complex distributions.

In this work, we continue this direction, but use *normalizing flows* [41, 37, 33] as our model class. The main advantage is that normalizing flows have successfully been able to represent complex data distributions in previous works and can be fit to such distributions in a straightforward manner. We thus rid ourselves almost entirely of the (potentially biased) constraints put previously on the inference models.

Summarizing, we show that:




Figure 1: Left: density plots of the inferred parameters  $q(\theta|\mathbf{x}_i)$  (some bold for visualization purposes) for each event  $\mathbf{x}_i$  in our simulated dataset. Right: despite the noise in the samples, the normalizing flow is able to recover the true population Gaussian mixture. We include the heuristic of averaging the noisy posterior samples, yielding an incorrect result.

- In the *Hierarchical Bayesian Analysis* (HBA) framework (a Bayesian graphical model used in the gravitational wave community), a normalizing flow can, despite having access to minimal and noisy data, recover structure in the data that agrees with either ground truth or robust results from previous studies. Further, it can reveal patterns that are easily missed by custom phenomenological modeling.
- Using a prior over the free parameters of the normalizing flow, we can get a posterior predictive distribution that generates interpretable confidence intervals.

## 2 Method

Let  $\mathbf{x} \in \mathbb{X} \subset \mathbb{R}^d$  be a gravitational wave time series and  $\mathbf{X} := \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$  be an observed dataset of such events. We understand the physical implication of an event  $\mathbf{x}$  through a vector of associated parameters  $\theta \in \mathbb{T} \subset \mathbb{R}^t$ . For binary mergers, these parameters include both masses, the redshift, and effective spin. Parameter inference for individual events comes with a lot of uncertainty due to measurement noise intrinsic to the underlying time series. Further, these posterior distributions can have a non-trivial shape. This is exemplified in fig. 1 and fig. 2, where we show on the left the density estimates of  $\theta_i$  associated with  $\mathbf{x}_i$  for every  $i = 1, \dots, N$ . It is clear that heuristics such as taking the mean of the samples of the posterior distributions and fitting the population model to those is not sufficient. Therefore, the gravitational wave community has resorted to *Hierarchical Bayesian Analysis* (HBA) [28, 16, 47], a Bayesian graphical model (see appendix A) for population-level inference. By taking into account measurement uncertainty, HBA distills the posterior samples into a single population model, which can be seen as a form of deconvolution [9]. In this way, despite huge uncertainties in the parameters of individual events, we can still draw scientifically sound downstream conclusions about important characteristics of the Universe.

### 2.1 Hierarchical Bayesian Analysis

We are interested in obtaining a population model for  $\theta$  after observing  $\mathbf{X}$ . First, we introduce a model parameter  $\mathbf{w} \in \mathbb{W} \subset \mathbb{R}^w$  with a fixed prior distribution  $p(\mathbf{w})$ . For a fully specified generative model, we further need to introduce distributions  $p(\theta_i|\mathbf{w})$  and  $p(\mathbf{x}_i|\theta_i)$ ,  $i = 1, \dots, N$ . We are interested in the *posterior distribution*  $p(\mathbf{w}|\mathbf{X}) \propto p(\mathbf{w}) \cdot \prod_{i=1}^N \int d\theta_i p(\theta_i|\mathbf{w}) p(\mathbf{x}_i|\theta_i)$

and the *posterior predictive distribution*

$$p(\theta_{N+1}|\mathbf{X}) = \int d\mathbf{w} p(\mathbf{w}|\mathbf{X}) p(\theta_{N+1}|\mathbf{w}) = \mathbb{E}_{p(\mathbf{w}|\mathbf{X})} [p(\theta_{N+1}|\mathbf{w})], \quad (1)$$

which expresses the probability of observing a new event with parameters  $\theta_{N+1}$  given the (old) observations  $\mathbf{X}$  while taking into account the uncertainty in  $\mathbf{w}$ .

While we usually do not have access to the marginal  $p(\theta_i)$  we can make use of a (usually physics-informed prior) distribution  $q(\theta_i)$ , and then receive samples from the conditional




Figure 2: Left: the GWTC-3 posterior density plots. Right: our inferred population model using normalizing flows for the inspiral *primary mass*  $p(m_1|\mathbf{X}) = \iint dm_1 dz d\mathcal{X}_{\text{eff}} p(\boldsymbol{\theta}|\mathbf{X})$  where we used a Monte Carlo estimate of eq. (1) to generate a 90% percentile interval.

$\boldsymbol{\theta}_i \sim q(\boldsymbol{\theta}_i|\mathbf{x}_i) := \frac{p(\mathbf{x}_i|\boldsymbol{\theta}_i)q(\boldsymbol{\theta}_i)}{q(\mathbf{x}_i)}$ , where  $q(\mathbf{x}_i) := \int d\boldsymbol{\theta}_i p(\mathbf{x}_i|\boldsymbol{\theta}_i)q(\boldsymbol{\theta}_i)$ . I.e., we assume that we have access to samples from the individual posterior distributions. As such, our dataset is  $\mathcal{D} := \{\{\boldsymbol{\theta}_{11}, \dots, \boldsymbol{\theta}_{1M_1}\}, \dots, \{\boldsymbol{\theta}_{N1}, \dots, \boldsymbol{\theta}_{NM_N}\}\}$ , where for every event  $i \in (1, \dots, N)$  we have  $M_i$  samples from  $q(\boldsymbol{\theta}_i|\mathbf{x}_i)$ . We are first interested in maximizing an approximate posterior using  $\mathcal{D}$ :

$$\mathbf{w}^* := \arg \max_{\mathbf{w} \in \mathbb{W}} \tilde{p}(\mathbf{w}|\mathbf{X}), \quad \tilde{p}(\mathbf{w}|\mathbf{X}) := \frac{p(\mathbf{w})}{Z(\mathbf{X}, \mathcal{D})} \prod_{i=1}^N \left( \frac{1}{M_i} \sum_{j=1}^{M_i} \left[ \frac{p(\boldsymbol{\theta}_{ij}|\mathbf{w})}{q(\boldsymbol{\theta}_{ij})} \right] \right), \quad (2)$$

with normalizing constant  $Z(\mathbf{X}, \mathcal{D})$  (derivation appendix B). A similar Monte-Carlo-based estimator is presented in [44]. After obtaining  $\mathbf{w}^*$  through, e.g., stochastic gradient descent, we run Markov chain Monte-Carlo around the MAP solution to obtain samples  $\mathbf{w} \sim \tilde{p}(\mathbf{w}|\mathbf{X})$ . These samples and our model  $p(\boldsymbol{\theta}|\mathbf{w})$  can be used to approximate the posterior predictive distribution eq. (1). It is important to model the whole posterior distribution, as the uncertainty in the parameters tells us where we can draw *robust* conclusions about the underlying population.

## 2.2 Normalizing Flows

So far, we have not discussed how we parameterize  $p(\boldsymbol{\theta}|\mathbf{w})$ . To facilitate increasingly larger datasets and to be able to model higher-dimensional distributions, the goal of this research is to explore more flexible and scalable models. As such, we propose *normalizing flows* [41, 37, 33, appendix C] as a suitable model class. Normalizing flows transform a density using a chain of a smooth, invertible mappings  $f: \mathbb{R}^t \rightarrow \mathbb{R}^t$  such that  $\boldsymbol{\theta}_K = f_K \circ \dots \circ f_2 \circ f_1(\boldsymbol{\theta}_0)$ , which constructs an arbitrarily complex density. We parameterize each  $f_k$  using a set of weights  $\mathbf{w}$ , and construct a log-likelihood  $\log p(\boldsymbol{\theta}_K|\mathbf{w}) = \log p_0(\boldsymbol{\theta}_0) - \sum_{k=1}^K \log \det \left| \frac{\partial f_k^{\mathbf{w}}}{\partial \boldsymbol{\theta}_k} \right|$ , which we can optimize for  $\mathbf{w}$  using, e.g., stochastic gradient descent. In this work, we use *planar flows* [37] and *block neural autoregressive flows* [12], as we experimentally found that they are sufficiently flexible while not using too many parameters. Furthermore, we found that the resulting distributions are *smooth*, which makes them physically more plausible and allows for better automated discovery of interpretable models [49].

## 3 Results

### 3.1 Illustrative Examples

In an illustrative example, we use a mixture of Gaussians as the ground truth marginal distribution over  $\boldsymbol{\theta}$ . The observable distribution is a Gaussian with mean  $|\boldsymbol{\theta}|$ , from which we draw  $\mathbf{X}$ . Using a uniform prior  $q(\boldsymbol{\theta})$ , we obtain posterior draws  $\boldsymbol{\theta} \sim q(\boldsymbol{\theta}|\mathbf{x}_i)$  for all events using rejection sampling. Due to the noninvertible forward model, the posteriors are bimodal and noisy. Since the true underlying population distribution and the posteriors are nontrivial, standard techniques will not be effective here. Still, the normalizing flow can, using the Hierarchical Bayesian Analysis framework, convert the noisy posterior draws shown in fig. 1 (left) into a faithful population model as shown in fig. 1 (right).

Further, we perform a more realistic (though less clearly visualizable) experiment with synthetic data in appendix D.2 where we sample data from previously suggested gravitational wave population models (e.g., by [4]) and recover the proposed distributions.

### 3.2 GWTC-3: Primary Mass Marginal

In our experiments on the GWTC-3 catalog, we consider four gravitational wave parameters:  $m_1$  (primary mass),  $m_2$  (secondary mass),  $z$  (redshift) and  $\mathcal{X}_{\text{eff}}$  (effective inspiral spin). I.e.,  $\theta := (m_1, m_2, z, \mathcal{X}_{\text{eff}})$  and  $t = 4$ . Note that the number of recorded gravitational wave events is currently rather small ( $N = 76$ ). Still, we show that we can retrieve structure from the data that agrees with previous studies. After getting  $p(\theta|\mathbf{w}^*)$  by maximizing eq. (2), we obtain, using Hamiltonian Monte Carlo [30],  $L = 10,000$  samples  $\mathbf{w} \sim p(\mathbf{w}|\theta)$ , starting the chain at  $\mathbf{w}^*$ . Next, we obtain the posterior predictive distribution by the approximation  $p(\theta|\mathbf{X}) \approx \mathbb{E}_{p(\mathbf{w}|\mathbf{X})}[p(\theta|\mathbf{w})] \approx \frac{1}{L} \sum_{l=1}^L p(\theta|\mathbf{w}_l)$ . We show the resulting *observed population* posterior predictive marginal  $p(m_1|\mathbf{X}) = \iiint dm_2 dz d\mathcal{X}_{\text{eff}} p(\theta|\mathbf{X}) = \iiint dm_2 dz d\mathcal{X}_{\text{eff}} p(m_1, m_2, z, \mathcal{X}_{\text{eff}}|\mathbf{X})$  in fig. 2.

Note, again, that we did not impose any strong restrictions on the model class; this result was fully recovered from just the data. Overall, we see (apart from selection bias effects [28]) similarity with, e.g., the POWER LAW + PEAK and MULTI-PEAK model of [5, 4]. Specifically, we observe distinct peaks around  $m_1 \sim 10 M_\odot$ ,  $m_1 \sim 40 M_\odot$ , and  $m_1 \sim 60 M_\odot$ . Finally, we also reassuringly observe that the uncertainty contracts or expands in the parts where we expect it to. I.e., it expands at regions where we do not have many observations, and contracts in the regions where we do (e.g., around the high-mass modes).

### 3.3 GWTC-3: Observed Population Distribution

We display the marginal joint distributions in fig. 3. Here, we define  $q := \frac{m_2}{m_1}$  as the ratio of secondary to primary mass of the binary. Note this figure shows the *observed population* results. This is especially noticeable in the redshift marginal  $z$ , which is heavily biased towards low-redshift events. Further, we expect more lower primary mass events  $m_1$ .  $\mathcal{X}_{\text{eff}}$  and  $q$  are expected to be less affected by not including the selection function. To draw conclusions about the true population, one should include *selection bias effects*. Optimization was highly unstable when including the selection function in the objective. Hence, including selection bias effects was left for future investigation. Also,  $q$  is expected to peak at  $q = 1$ , with an infinitely steep cutoff at  $q > 1$ . We did not observe this due to the smoothness of the flow. Future work could look into reparameterizations to enable this behavior.

We can tentatively draw a few conclusions from the observed population figure. Interestingly, it suggests that  $m_1$  and  $q$  do not simply correlate through a power law. Instead, the distribution shows an inverted V shape. The redshift to mass correlations seem plausible and consistent. The effective spin marginal also aligns with previous literature. In [29] it was noted that  $q$  does not correlate positively with  $\mathcal{X}_{\text{eff}}$ , which is also confirmed here.

## 4 Discussion

We parameterized the population model in the Hierarchical Bayesian Analysis framework with a normalizing flow. In illustrative experiments, we confirmed that the model can retrieve the actual population density under nontrivial posterior samples. On gravitational wave data from GWTC-3, we considered primary mass, secondary mass, redshift, and effective spin and recovered an observed

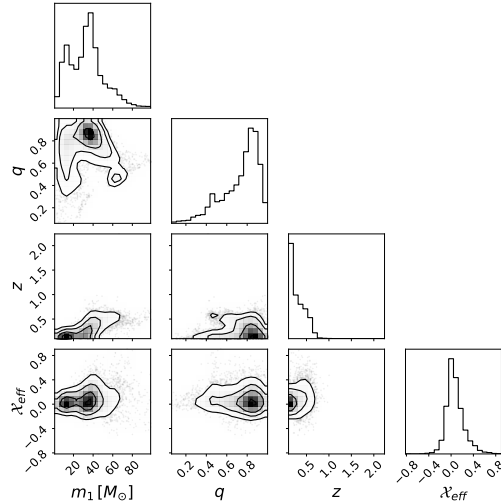


Figure 3: Pair plot of joint marginals of our inferred  $p(\theta|\mathbf{X})$ . Note that we did not take in to account selection bias effects which is especially noticeable in the redshift marginal. Hence, this is an *observed population* analysis.

population model that agrees with robust previous phenomenological modeling results despite the dataset being small and highly noisy. This paves the way for less constrained, tractable, and flexible population-level inference in noisy settings, especially once more higher-dimensional data become available, allowing for automated discovery of structure in the data.

## 5 Broader Impact

This research proposed to combine normalizing flows with the HBA framework. The current work found an application to gravitational wave event data, but in principle, the method can be implemented in many fields of science. While this work purely aims to aid scientific discovery, the approach can also be used in non-scientific settings dealing with noisy observations, which can pose privacy-related issues depending on the data type and source. As such, we encourage more data privacy awareness and establishing data protection policies.

## 6 Acknowledgements

This work was largely performed as part of the machine learning summer program of the Flatiron Institute. We would like to thank the scientific software development community, without whom this work would not be possible [20, 43, 36, 7, 35, 46, 11, 48].

## References

- [1] BP Abbott, Richard Abbott, TDea Abbott, S Abraham, F Acernese, K Ackley, C Adams, RX Adhikari, VB Adya, Christoph Affeldt, et al. Gwtc-1: a gravitational-wave transient catalog of compact binary mergers observed by ligo and virgo during the first and second observing runs. *Physical Review X*, 9(3):031040, 2019.
- [2] R Abbott, TD Abbott, S Abraham, F Acernese, K Ackley, A Adams, C Adams, RX Adhikari, VB Adya, Christoph Affeldt, et al. Gwtc-2: compact binary coalescences observed by ligo and virgo during the first half of the third observing run. *Physical Review X*, 11(2):021053, 2021.
- [3] R Abbott, TD Abbott, F Acernese, K Ackley, C Adams, N Adhikari, RX Adhikari, VB Adya, C Affeldt, D Agarwal, et al. Gwtc-3: compact binary coalescences observed by ligo and virgo during the second part of the third observing run. *arXiv preprint arXiv:2111.03606*, 2021.
- [4] R Abbott, TD Abbott, F Acernese, K Ackley, C Adams, N Adhikari, RX Adhikari, VB Adya, C Affeldt, D Agarwal, et al. The population of merging compact binaries inferred using gravitational waves through gwtc-3. *arXiv preprint arXiv:2111.03634*, 2021.
- [5] Rich Abbott, TD Abbott, S Abraham, F Acernese, K Ackley, A Adams, C Adams, RX Adhikari, VB Adya, Christoph Affeldt, et al. Population properties of compact objects from the second ligo–virgo gravitational-wave transient catalog. *The Astrophysical journal letters*, 913(1):L7, 2021.
- [6] Rianne van den Berg, Leonard Hasenclever, Jakub M Tomczak, and Max Welling. Sylvester normalizing flows for variational inference. *arXiv preprint arXiv:1803.05649*, 2018.
- [7] Eli Bingham, Jonathan P Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul Szerlip, Paul Horsfall, and Noah D Goodman. Pyro: Deep universal probabilistic programming. *The Journal of Machine Learning Research*, 20(1):973–978, 2019.
- [8] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International conference on machine learning*, pages 1613–1622. PMLR, 2015.
- [9] Jo Bovy, David W Hogg, and Sam T Roweis. Extreme deconvolution: Inferring complete distribution functions from noisy, heterogeneous and incomplete observations. *The Annals of Applied Statistics*, 5(2B):1657–1677, 2011.
- [10] Thomas A Callister, Carl-Johan Haster, Ken KY Ng, Salvatore Vitale, and Will M Farr. Who ordered that? unequal-mass binary black hole mergers have larger effective spins. *The Astrophysical Journal Letters*, 922(1):L5, 2021.

- [11] Adam D Cobb and Brian Jalaian. Scaling hamiltonian monte carlo inference for bayesian neural networks with symmetric splitting. *Uncertainty in Artificial Intelligence*, 2021.
- [12] Nicola De Cao, Wilker Aziz, and Ivan Titov. Block neural autoregressive flow. In *Uncertainty in artificial intelligence*, pages 1263–1273. PMLR, 2020.
- [13] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- [14] Maya Fishbach, Daniel E Holz, and Will M Farr. Does the black hole merger rate evolve with redshift? *The Astrophysical Journal Letters*, 863(2):L41, 2018.
- [15] Maya Fishbach, Daniel E. Holz, and Will M. Farr. Does the Black Hole Merger Rate Evolve with Redshift? *Astrophys. J. Lett.*, 863(2):L41, 2018.
- [16] Sebastian M Gaebel, John Veitch, Thomas Dent, and Will M Farr. Digging the population of compact binary mergers out of the noise. *Monthly Notices of the Royal Astronomical Society*, 484(3):4008–4023, 2019.
- [17] Yarín Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016.
- [18] Shanika Galaudage, Colm Talbot, and Eric Thrane. Gravitational-wave inference in the catalog era: Evolving priors and marginal events. *Physical Review D*, 102(8):083026, 2020.
- [19] Emiel Hoogeboom, Victor Garcia Satorras, Jakub Tomczak, and Max Welling. The convolution exponential and generalized sylvester flows. *Advances in Neural Information Processing Systems*, 33:18249–18260, 2020.
- [20] John D Hunter. Matplotlib: A 2d graphics environment. *Computing in science & engineering*, 9(03):90–95, 2007.
- [21] Shasvath J Kapadia, Sarah Caudill, Jolien DE Creighton, Will M Farr, Gregory Mendell, Alan Weinstein, Kipp Cannon, Heather Fong, Patrick Godwin, Rico KL Lo, et al. A self-consistent method to estimate the rate of compact binary coalescences with a poisson mixture model. *Classical and Quantum Gravity*, 37(4):045007, 2020.
- [22] Chase Kimball, Colm Talbot, Christopher PL Berry, Matthew Carney, Michael Zevin, Eric Thrane, and Vicky Kalogera. Black hole genealogy: Identifying hierarchical mergers with gravitational waves. *The Astrophysical Journal*, 900(2):177, 2020.
- [23] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [24] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018.
- [25] Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. *Advances in neural information processing systems*, 29, 2016.
- [26] Ely D Kovetz, Ilias Cholis, Patrick C Breysse, and Marc Kamionkowski. Black hole mass function from gravitational wave measurements. *Physical Review D*, 95(10):103010, 2017.
- [27] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
- [28] Ilya Mandel, Will M Farr, and Jonathan R Gair. Extracting distribution parameters from multiple uncertain observations with selection biases. *Monthly Notices of the Royal Astronomical Society*, 486(1):1086–1093, 2019.
- [29] B McKernan, KES Ford, T Callister, WM Farr, R O’Shaughnessy, R Smith, E Thrane, and A Vajpeyi. Ligo–virgo correlations between mass ratio and effective inspiral spin: testing the active galactic nuclei channel. *Monthly Notices of the Royal Astronomical Society*, 514(3):3886–3893, 2022.
- [30] Radford M Neal et al. Mcmc using hamiltonian dynamics. *Handbook of markov chain monte carlo*, 2(11):2, 2011.

- [31] Alexander H Nitz, Collin Capano, Alex B Nielsen, Steven Reyes, Rebecca White, Duncan A Brown, and Badri Krishnan. 1-ogc: The first open gravitational-wave catalog of binary mergers from analysis of public advanced ligo data. *The Astrophysical Journal*, 872(2):195, 2019.
- [32] Alexander H Nitz, Thomas Dent, Gareth S Davies, Sumit Kumar, Collin D Capano, Ian Harry, Simone Mozzon, Laura Nuttall, Andrew Lundgren, and Márton Tápai. 2-ogc: Open gravitational-wave catalog of binary mergers from analysis of public advanced ligo and virgo data. *The Astrophysical Journal*, 891(2):123, 2020.
- [33] George Papamakarios, Eric T Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *J. Mach. Learn. Res.*, 22(57):1–64, 2021.
- [34] George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. *Advances in neural information processing systems*, 30, 2017.
- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- [36] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
- [37] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR, 2015.
- [38] Javier Roulet, Tejaswi Venumadhav, Barak Zackay, Liang Dai, and Matias Zaldarriaga. Binary black hole mergers from ligo/virgo o1 and o2: Population inference combining confident and marginal events. *Physical Review D*, 102(12):123022, 2020.
- [39] Javier Roulet and Matias Zaldarriaga. Constraints on binary black hole populations from ligo–virgo detections. *Monthly Notices of the Royal Astronomical Society*, 484(3):4216–4229, 2019.
- [40] Rory JE Smith, Colm Talbot, Francisco Hernandez Vivanco, and Eric Thrane. Inferring the population properties of binary black holes from unresolved gravitational waves. *Monthly Notices of the Royal Astronomical Society*, 496(3):3281–3290, 2020.
- [41] Esteban G Tabak and Eric Vanden-Eijnden. Density estimation by dual ascent of the log-likelihood. *Communications in Mathematical Sciences*, 8(1):217–233, 2010.
- [42] Vaibhav Tiwari. Vamana: modeling binary black hole population with minimal assumptions. *Classical and Quantum Gravity*, 38(15):155007, 2021.
- [43] Stefan Van Der Walt, S Chris Colbert, and Gael Varoquaux. The numpy array: a structure for efficient numerical computation. *Computing in science & engineering*, 13(2):22–30, 2011.
- [44] Maxime Vandegar, Michael Kagan, Antoine Wehenkel, and Gilles Louppe. Neural empirical bayes: Source distribution estimation and its applications to simulation-based inference. In *International Conference on Artificial Intelligence and Statistics*, pages 2107–2115. PMLR, 2021.
- [45] Tejaswi Venumadhav, Barak Zackay, Javier Roulet, Liang Dai, and Matias Zaldarriaga. New binary black hole mergers in the second observing run of advanced ligo and advanced virgo. *Physical Review D*, 101(8):083030, 2020.
- [46] Pauli Virtanen, Ralf Gommers, Travis E Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, et al. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nature methods*, 17(3):261–272, 2020.
- [47] Salvatore Vitale, Davide Gerosa, Will M Farr, and Stephen R Taylor. Inferring the properties of a population of compact binaries in presence of selection effects. In *Handbook of Gravitational Wave Astronomy*, pages 1–60. Springer, 2022.
- [48] Michael L Waskom. Seaborn: statistical data visualization. *Journal of Open Source Software*, 6(60):3021, 2021.

- [49] Kaze W. K Wong and Miles Cranmer. Automated discovery of interpretable gravitational-wave population models, 2022.
- [50] Daniel Wysocki, Jacob Lange, and Richard O’Shaughnessy. Reconstructing phenomenological distributions of compact binaries via gravitational wave observations. *Physical Review D*, 100(4):043012, 2019.
- [51] Barak Zackay, Tejaswi Venumadhav, Liang Dai, Javier Roulet, and Matias Zaldarriaga. Highly spinning and aligned binary black hole merger in the advanced ligo first observing run. *Physical Review D*, 100(2):023007, 2019.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
  - (b) Did you describe the limitations of your work? [Yes] E.g., the exclusion of selection bias effects.
  - (c) Did you discuss any potential negative societal impacts of your work? [Yes] In the broader impact statement.
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
  - (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Code will be released upon acceptance.
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] In the supplementary material.
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? [Yes] Will be done in acknowledgements statement.
  - (b) Did you mention the license of the assets? [N/A]
  - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
  - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A]
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
  - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]






Figure 4: Graphical model of Hierarchical Bayesian Analysis. The red arrows are the ones that our normalizing flow parameterizes.

## A Graphical Model

The graphical we use is shown in fig. 4. We colored the arrows that our normalizing flow parameterizes red. To estimate  $p(\theta_{N+1} | \mathbf{X})$ , samples  $\theta_i \sim q(\theta_i | \mathbf{x}_i)$  are inferred for all  $\mathbf{x}_i$  in the first stage. In this work we assume we have those samples and use them to estimate the remaining densities in eq. (1) through the objective eq. (2).

## B Hierarchical Bayesian Likelihood

Using  $q(\theta_i | \mathbf{x}_i) := \frac{p(\mathbf{x}_i | \theta_i) q(\theta_i)}{q(\mathbf{x}_i)}$ , we set

$$\hat{p}(\mathbf{x} | \mathbf{w}) := \int d\theta \left[ q(\theta | \mathbf{x}) \cdot \frac{q(\mathbf{x})}{q(\theta)} \right] p(\theta | \mathbf{w}) \quad (3)$$

$$= q(\mathbf{x}) \cdot \mathbb{E}_{q(\theta | \mathbf{x})} \left[ \frac{p(\theta | \mathbf{w})}{q(\theta)} \right]. \quad (4)$$

Note that  $q(\mathbf{x}_i)$  is a constant with respect to our model parameters  $\mathbf{w}$ . Hence,

$$\hat{p}(\mathbf{w} | \mathbf{X}) := \frac{p(\mathbf{w})}{Z(\mathbf{X}, \mathcal{D})} \prod_{i=1}^N \mathbb{E}_{q(\theta | \mathbf{x}_i)} \left[ \frac{p(\theta | \mathbf{w})}{q(\theta)} \right] \quad (5)$$

$$\approx \frac{p(\mathbf{w})}{Z(\mathbf{X}, \mathcal{D})} \prod_{i=1}^N \left( \frac{1}{M_i} \sum_{j=1}^{M_i} \left[ \frac{p(\theta_j | \mathbf{w})}{q(\theta_j)} \right] \right) =: \tilde{p}(\mathbf{w} | \mathbf{X}) \quad (6)$$

## C Normalizing Flows

To facilitate increasingly larger datasets and to be able to model higher-dimensional distributions, the goal of this research is to explore more flexible and scalable models. As such, we propose *normalizing flows*[41, 37, 33] as a suitable model class. Normalizing flows transform a density using a smooth, invertible mapping  $f : \mathbb{R}^t \rightarrow \mathbb{R}^t$  with inverse  $g := f^{-1}$ . As such, if we transform a random variable  $\theta$ , then  $\theta'$  has

$$p(\theta') = p(\theta) \left| \det \frac{\partial f^{-1}}{\partial \theta'} \right|. \quad (7)$$

If we use a chain of these transformations, such that

$$\theta_K = f_K \circ \dots \circ f_2 \circ f_1(\theta_0), \quad (8)$$

we can construct an arbitrarily complex density. We parameterize each  $f_k$  using a set of weights  $\mathbf{w}$ , and construct a log-likelihood

$$\log p(\theta_K | \mathbf{w}) = \log p_0(\theta_0) - \sum_{k=1}^K \log \det \left| \frac{\partial f_k^{\mathbf{w}}}{\partial \theta_k} \right|, \quad (9)$$

which we can optimize for  $\mathbf{w}$  using, e.g., stochastic gradient descent. The downside is a potentially computationally expensive Jacobian determinant. As such, much work has been conducted on finding transformations that have computationally cheap Jacobian determinants [13, 25, 34, 6, 24, 19]. In this work, we use *planar flows* [37] and *block neural autoregressive flows* [12], as we experimentally found that they are sufficiently flexible while not using too many parameters. Furthermore, we found that the resulting distributions are (relative to some other parameterizations) *smooth*, which makes them physically more plausible.

**Planar Flow** Planar flow uses the transformation

$$f_k^{\mathbf{w}}(\boldsymbol{\theta}) := \boldsymbol{\theta} + \mathbf{u}_k h(\mathbf{v}_k^\top \boldsymbol{\theta} + b_k), \quad (10)$$

where  $(\mathbf{u}_k, \mathbf{z}_k, b_k) \in \mathbf{w}$  are free parameters, and  $h$  is the hyperbolic tangent. These transformations contract and expand the inputs perpendicular to the plane defined  $\mathbf{v}_k^\top \boldsymbol{\theta} - b_k = 0$ . It can be shown that the Jacobian determinant of this mapping can be computed in  $O(t)$  [37]. As such, we can use a sequence of these transformations to model  $p(\boldsymbol{\theta}|\mathbf{w})$ , our gravitational wave population model (eq. (2)).

**Block Neural Autoregressive Flow** Neural autoregressive flows decomposes a joint distribution over  $\boldsymbol{\theta} \in \mathbb{R}^t$  into  $t$  conditional distributions. the transformations  $f_k^{\mathbf{w}}$  then yield lower triangular Jacobians and hence cheaply computable determinants. Block neural autoregressive flows then parameterize  $f_k^{\mathbf{w}}$  using a neural network that uses block lower-triangular matrices. Their diagonal elements are strictly positive, ensuring monotonicity. The off-diagonal elements do not require monotonicity as they only play a role in the conditioning part of the conditional distributions.

## D Experimental Details

All experiments were conducted locally.

### D.1 Illustrative Example

The ground-truth population is defined by a mixture  $\boldsymbol{\theta} \sim \mathcal{N}(-0.75, 0.5^2)\mathcal{N}(0.75, 0.5^2)$ . The forward model is given by  $\mathbf{x}|\boldsymbol{\theta} \sim N(|\boldsymbol{\theta}|, 0.1^2)$ . This noninvertible forward mapping results in noisy, bimodal posteriors. We sample  $N = 1024$  events. Using a uniform prior  $q(\boldsymbol{\theta}) := U(-5, 5)$  and the true forward model, we draw  $M_i = 128$  posterior samples using rejection sampling from  $q(\boldsymbol{\theta}|\mathbf{x}_i)$  for every  $\mathbf{x}_i, i \in 1, \dots, 1024$ . We parameterize  $p(\boldsymbol{\theta}|\mathbf{w})$  as a 4-layer planar flow and a standard Gaussian base distribution. We use the Adam optimizer [23] with learning rate 0.01 so obtain  $\mathbf{w}^*$ . We found empirically that optimization runs better when mini-batches are taken for both the events and also the posterior draws. In this experiment, we batch sizes of  $N' := 1024$  events and  $M'_i := 4$ .

### D.2 Synthetic Data Experiment

Here, we simulate a population model using previously published models and investigate whether our flow-based approach can faithfully recover those distributions. We start with the POWER LAW + PEAK model for  $p(m_1, m_2)$  from [4] using the means of the posterior distributions for ground-truth population parameters. For a conditional effective spin distribution we use [10], again with the mean of the inferred distributions as ground-truth parameter values. Finally, we model the distribution of redshift using a power law with spectral index  $\kappa = 1.7$  following [15]. Note that there is an extra factor in the merger coming from the comoving volume factor. The resulting samples are shown in fig. 5 (left). We simulate Gaussian posteriors in our Hierarchical Bayesian framework. The distribution that the normalizing flow recovers is shown on the right in fig. 5. We see that the flow recovers most of the correlations and shapes of the joint distribution faithfully.

We reparameterize the  $m_1$  samples to log-space and  $q$  samples to logit-space. We sample  $N = 1024$  events and use 32 hierarchical samples per batch. We used the Adam optimizer [23] with learning rate of 0.001 and 1 layer of block neural autoregressive flow with 2 8-dimensional hidden layers.




Figure 5: Left: simulated gravitational wave data. Right: recovered distributions by the normalizing flow.

### D.3 GWTC-3

We have  $N = 76$  gravitational wave events available. The number of posterior samples per event ranges from  $M_i = 3194$  to  $M_i = 268806$ . We use 10 layers of *planar flow*, resulting in a total of  $w = 90$  free parameters, limiting the risk of severe over-fitting resulting in unnatural population distribution shapes. We use mini-batches of  $N' := 76$  events and  $M'_i := 32$  posterior samples per event. We use the Adam optimizer [23] with learning rate 0.001. to find  $\mathbf{w}^*$ .

## E Future Work

In this study, we did not include *selection bias* effects. Future work should adjust for this to obtain more realistic models of population parameters of gravitational waves. We also expect that we did not fully explore all parameterizations (and hyperparameters) that recover the population model as faithfully as possible, both in terms of the normalizing flow and how the input data should be transformed for the flow to optimally learn the densities.

Further, [49] have presented a promising approach to distill black-box models into symbolic equations. Such an approach can also be applied to our normalizing flows, yielding more interpretable versions of the presented population densities.

Finally, we used Hamiltonian Monte-Carlo to get samples  $\mathbf{w} \sim p(\mathbf{w}|\boldsymbol{\theta})$ . This is not scalable when the number of weights increases in order to be able to fit more complex data. To resolve this, follow-up works could explore approaches based on the *Bayesian Neural Network* literature, e.g., [8, 17, 27].