

Gravitational-wave parameter estimation with autoregressive neural network flows

Stephen R. Green^{1,*}, Christine Simpson^{2,†} and Jonathan Gair^{1,‡}

¹Max Planck Institute for Gravitational Physics (Albert Einstein Institute) Am Mühlenberg 1, 14476 Potsdam, Germany

²School of Informatics, University of Edinburgh, 10 Crichton Street, Edinburgh EH8 9AB, United Kingdom



(Received 26 February 2020; accepted 28 September 2020; published 19 November 2020)

We introduce the use of autoregressive normalizing flows for rapid likelihood-free inference of binary black hole system parameters from gravitational-wave data with deep neural networks. A normalizing flow is an invertible mapping on a sample space that can be used to induce a transformation from a simple probability distribution to a more complex one: if the simple distribution can be rapidly sampled and its density evaluated, then so can the complex distribution. Our first application to gravitational waves uses an autoregressive flow, conditioned on detector strain data, to map a multivariate standard normal distribution into the posterior distribution over system parameters. We train the model on artificial strain data using a model for the gravitational-wave signal that includes inspiral, merger and ringdown and draw waveforms from a five-parameter $(m_1, m_2, \phi_0, t_c, d_L)$ prior and stationary Gaussian noise realizations with a fixed power spectral density. This gives performance comparable to current best deep-learning approaches to gravitational-wave parameter estimation. We then build a more powerful latent variable model by incorporating autoregressive flows within the variational autoencoder framework. This model has performance comparable to Markov chain Monte Carlo and, in particular, successfully models the multimodal ϕ_0 posterior. Finally, we train the autoregressive latent variable model on an expanded parameter space, including also aligned spins (χ_{1z}, χ_{2z}) and binary inclination θ_{JN} , and show that all parameters and degeneracies and most uncertainties are well-recovered. In all cases, sampling is extremely fast, requiring less than two seconds to draw 10^4 posterior samples.

DOI: [10.1103/PhysRevD.102.104057](https://doi.org/10.1103/PhysRevD.102.104057)

I. INTRODUCTION

The task of gravitational-wave parameter estimation is to determine the system parameters that gave rise to observed detector strain data. This is accomplished using Bayesian inference. Assuming a *likelihood* model $p(y|x)$ for strain data y conditioned on system parameters x , and a *prior* distribution $p(x)$, one obtains through Bayes' theorem the *posterior* distribution,

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

*stephen.green@aei.mpg.de

†christine.simpson@ed.ac.uk

‡jonathan.gair@aei.mpg.de

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Open access publication funded by the Max Planck Society.

where the normalization $p(y)$ is called the model *evidence*. Generally, one can evaluate the likelihood explicitly (although this may be computationally expensive) and the prior is also known, so this allows for calculation of the posterior up to normalization. One can then use an algorithm such as Markov chain Monte Carlo (MCMC) to obtain samples from the posterior.

Standard sampling algorithms are effective, but they can be computationally costly. Indeed, for binary black holes, obtaining a sufficient number of posterior samples can take on the order of days, whereas for binary neutron stars, this extends to weeks [1,2]. Especially for binary neutron stars, which might have multimessenger counterparts [3], it is critical to reduce this time to provide accurate information to electromagnetic observers.

There have recently been several efforts to speed up parameter estimation by using deep learning [4–6]. The main tool of deep learning is the neural network, which is a trainable and very flexible function approximator. Neural networks can have millions of parameters, which are tuned through stochastic gradient descent to optimize a loss function. In this way, very complex functions can be

approximated. Since conditional distributions can be parametrized by functions, neural networks can be used to model gravitational-wave posteriors.

A key observation of [4,5] is that Bayes' rule can be applied in such a way that training only requires samples from the prior and the gravitational-wave likelihood, i.e., strain realizations $y \sim p(y|x)$. Although the network learns the posterior distribution, posterior samples are not needed for training. Training also does not require any likelihood evaluations, and for this reason this approach is known as *likelihood-free inference*. Since obtaining posterior samples and evaluating the likelihood are expensive, the likelihood-free approach is very fast in comparison. It is also applicable in contexts where a simulation of the data generative process is available, but the likelihood function may be unknown.

In [4], this approach was applied successfully with a multivariate Gaussian posterior model, i.e., $p(x|y) = \mathcal{N}(\mu(y), \Sigma(y))$, where the mean $\mu(y)$ and covariance matrix $\Sigma(y)$ are given by neural networks. Once trained on simulated waveforms and noise realizations, the Gaussian posterior model is trivial to sample.

The challenge, then, is to define a sufficiently flexible model distribution for the posterior. Gaussians are good approximations for very high signal-to-noise ratio, but posteriors can in general have higher moments and multimodality. Reference [4] also experimented with Gaussian mixture models and posterior histogram outputs. The performance of the Gaussian mixture, however, was not a significant improvement over the single Gaussian, and although the histogram was effective, it is limited to describing one or two parameters.

A promising approach to increase the flexibility of a posterior model is to introduce latent variables z and perform variational Bayesian inference. The posterior is then given by marginalization over the latent variables, i.e., $p(x|y) = \int p(x|z, y)p(z|y)dz$. With $p(x|z, y)$ and $p(z|y)$ both taken to be Gaussian, this nevertheless gives *non-Gaussian* $p(x|y)$. This can be implemented using the variational autoencoder framework [7,8], and recent results [5] for gravitational-wave parameter estimation with a conditional variational autoencoder (CVAE) have demonstrated the recovery of non-Gaussian posteriors (although not multimodality).

In this work, we use the method of *normalizing flows* [9] to define very flexible model distributions with improved ability to approximate complex gravitational-wave posteriors. A normalizing flow $f: X \rightarrow X$ is an invertible mapping on a sample space X , with a simple Jacobian determinant. Using the change of variables rule for probabilities, this induces a transformation,

$$p(x) = \pi(f^{-1}(x)) \left| \det \frac{\partial(f_1^{-1}, \dots, f_n^{-1})}{\partial(x_1, \dots, x_n)} \right|, \quad (2)$$

from a base distribution $\pi(u)$ to a new distribution $p(x)$. Here $n = \dim X$. Starting from a simple standard normal distribution for $\pi(u)$, one can obtain a more complex distribution by applying a normalizing flow. To describe a conditional distribution, such as a gravitational-wave posterior, a normalizing flow can be made to depend on additional variables; for our application, we take $f = f_y(u)$ to depend on detector strain data y .

The fact that f is invertible means that the normalizing flow enables both sampling and density evaluation of $p(x)$, provided this is possible for $\pi(u)$. To sample $x \sim p(x)$, one first samples $u \sim \pi(u)$ and then sets $x = f(u)$. For given x , to evaluate the density, the inverse mapping $u = f^{-1}(x)$ is used in the right-hand side of (2).

Normalizing flows may be used to model gravitational-wave posterior distributions directly, and this is the first application that we describe. We fix the base distribution $\pi(u)$ to be a multivariate standard normal of the same dimension as the system parameter space. We define the normalizing flow $f_y(u)$ in terms of a neural network, as a *masked autoregressive flow (MAF)* [10]. The posterior model $p(x|y)$ is then defined through (2). Since the density can be evaluated directly and differentiated with respect to the neural network parameters defining the flow, the model can be trained using stochastic gradient descent to maximize the likelihood that the training data $[(x, y) \text{ pairs}]$ came from the model.

Normalizing flows can also be incorporated into the CVAE: they can be used to enhance the flexibility of the encoder, decoder, and prior component networks, thereby increasing the overall flexibility of the model [11,12]. In Sec. II of this work, we describe all of the above networks in further detail.

In Sec. III we present the results of our experiments in using these models to describe gravitational-wave posteriors over a five-dimensional space of system parameters. Following [5], we study binary black hole waveforms over the parameters $(m_1, m_2, \phi_0, t_c, d_L)$, with added noise drawn from a stationary Gaussian distribution with fixed power spectral density. We work in the time domain. We find that the basic MAF network achieves results comparable to [5], but with the advantage of not having any latent variables to marginalize over. In our experiments, however, neither of these networks successfully models the posterior over ϕ_0 , which is multimodal. We next test the more powerful autoregressive CVAE, which does succeed in modeling the multimodality in ϕ_0 . To validate our recovered posteriors, we present p-p plots consistent with uniformly distributed percentile scores in each parameter, as well as comparisons to MCMC sampling.

We then expand the parameter space to include aligned spins (χ_{1z}, χ_{2z}) and binary inclination θ_{JN} in Sec. IV. We train the CVAE network with autoregressive flows to model the posterior over this eight-dimensional space. We find that the network once again successfully models all

TABLE I. Vector spaces.

Space	Dimension	Description
X	n	System parameters
Y	m	Strain data
Z	l	Latent variables

parameters, although the mass uncertainty is slightly overestimated. The various degeneracies, e.g., between χ_{1z} and χ_{2z} , and between θ_{JN} and d_L , are well recovered. We validate our results again with a p-p plot.

This work is organized as follows. In the following section we describe in more detail the various types of neural networks that we use for parameter estimation. In Sec. III we describe our experiments with the five-dimensional parameter space, and in Sec. IV the enlarged eight-dimensional space. Finally, in Sec. V we conclude with a discussion of potential further improvements.

Software: All of our neural networks are implemented in PyTorch [13], with the autoregressive network implemented using Pyro [14]. We used emcee [15] to generate MCMC comparisons and ChainConsumer [16] to produce corner plots.

Notation: The various vector spaces, along with their dimensionalities are given in Table I.

II. PRELIMINARIES

In this section we review important deep-learning concepts, and we discuss them in the context of gravitational-wave parameter estimation. The first two subsections describe ideas that have already been implemented for parameter estimation, in [4,5], respectively, and the last two describe the autoregressive flows that we explore in this work.

A. Neural network models of gravitational-wave posteriors

Suppose we have a posterior distribution $p_{\text{true}}(x|y)$. Our aim is to train a neural network to give an approximation $p(x|y)$ to $p_{\text{true}}(x|y)$. The “true” posterior is itself a model for the physical system. For gravitational waves, it is defined through Bayes’ theorem in terms of a prior $p_{\text{true}}(x)$ over the system parameters x and a likelihood $p_{\text{true}}(y|x)$. The likelihood depends on a choice of waveform model h and a measured noise power spectral density (PSD) $S_n(f)$: it is the probability density that the residual $n = y - h(x)$ is drawn from a stationary Gaussian noise distribution with PSD $S_n(f)$. For further details on the noise model and likelihood, see, e.g., [17].

For stationary Gaussian noise and known $h(x)$, it is trivial to sample from the likelihood. In contrast, the “inverse problem” of sampling from the posterior—sampling over parameters x —requires an algorithm such as MCMC and many evaluations of the waveform model

and the likelihood. This repeated comparison between model waveforms and data is computationally expensive, and for this reason we wish to develop the neural network model, $p(x|y)$.

The first step in developing the model is to parametrize the posterior in terms of a neural network. For now, we take as our model a multivariate normal distribution [4], although our main interest later will be in defining more flexible models. That is, we take

$$p(x|y) = \frac{1}{\sqrt{(2\pi)^n |\det \Sigma(y)|}} \times \exp \left(-\frac{1}{2} \sum_{ij=1}^n (x_i - \mu_i(y)) \Sigma_{ij}^{-1}(y) (x_j - \mu_j(y)) \right), \quad (3)$$

where the mean $\mu(y)$ and covariance matrix $\Sigma(y)$ are functions of data y defined by a feedforward neural network. [To ensure $\Sigma(y)$ is positive definite and symmetric, it is in practice better to take the Cholesky decomposition, $\Sigma(y) = A(y)A(y)^\top$, where $A(y)$ is lower triangular with positive diagonal entries. $A(y)$ is then modeled with the neural network.]

Feedforward neural networks can have a variety of configurations, but the simplest consists of a sequence of fully connected layers. The output of the first *hidden* layer (of dimension d_1) is

$$h_1 = \sigma(W_1 y + b_1), \quad (4)$$

where W_1 is a $d_1 \times m$ matrix, and b_1 is a d_1 -dimensional vector. σ is an element-wise nonlinearity, typically taken to be a *rectified linear unit*,

$$\sigma(u) = \begin{cases} u & \text{if } u > 0, \\ 0 & \text{if } u \leq 0. \end{cases} \quad (5)$$

The output h_1 is then passed through a second hidden layer,

$$h_2 = \sigma(W_2 h_1 + b_2), \quad (6)$$

and so on, for as many hidden layers as desired. A final affine transformation $W_p h_{p-1} + b_p$ is then applied, and the outputs repackaged into a vector μ and lower-triangular matrix A . A suitable nonlinearity should be applied to ensure positive diagonal components of A , but otherwise the components of μ and A should be unconstrained. The weight matrices W_i and bias vectors b_i are initialized randomly (see, e.g., [18]) and then trained through stochastic gradient descent to optimize a loss function.

With the posterior defined, the loss function should be chosen so that after training, $p(x|y)$ is a good approximation to $p_{\text{true}}(x|y)$. We therefore take it to be the expectation value (over y) of the cross entropy between the two distributions,

$$L = - \int dx dy p_{\text{true}}(x|y) p_{\text{true}}(y) \log p(x|y), \quad (7)$$

i.e., we maximize the likelihood over the network parameters (the weights and biases) that the training data came from the model. Minimizing L is equivalent to minimizing the expectation value of the Kullback-Leibler (KL) divergence between the true and model posteriors, since $p_{\text{true}}(x|y)$ is fixed.

The loss function in the form (7) is actually not ideal for our purposes. The reason is that it requires sampling from $p_{\text{true}}(x|y)$ —a very costly operation. Instead, as pointed out by [4,5], we can use Bayes' theorem in a very advantageous way, to write

$$L = - \int dx dy p_{\text{true}}(y|x) p_{\text{true}}(x) \log p(x|y). \quad (8)$$

To train the network now requires sampling from the likelihood, not the posterior. On a minibatch of training data of size N , we approximate

$$L \approx - \frac{1}{N} \sum_{i=1}^N \log p(x^{(i)}|y^{(i)}), \quad (9)$$

where $x^{(i)} \sim p_{\text{true}}(x)$ and $y^{(i)} \sim p_{\text{true}}(y|x^{(i)})$. The loss function can be explicitly evaluated using expression (3).

The gradient of L with respect to the network parameters can be calculated using backpropagation (i.e., the chain rule), and the network optimized with stochastic gradient descent. The training set consists of parameter samples x and waveforms $h(x)$; random noise realizations are added at train time to obtain data samples y .

B. Variational autoencoders

One way to increase the flexibility of the model is to introduce latent variables z , and define the gravitational-wave posterior by first sampling from a prior over z , and then from a distribution over x , conditional on z . In other words,

$$p(x|y) = \int dz p(x|z, y) p(z|y). \quad (10)$$

If one takes $p(x|z, y)$ and $p(z|y)$ to both be multivariate Gaussians, then $p(x|y)$ is a Gaussian mixture of Gaussians. In this way one can describe a more flexible posterior.

At first glance it is not clear how to train such a model: the posterior (10) is intractable, since evaluation involves marginalizing over z . If one knew the posterior¹ $p(z|x, y)$, then

¹The variational posterior $p(z|x, y)$ should not be confused with the gravitational-wave posterior $p(x|y)$. It should be clear from context to which posterior distribution we are referring.

$$p(x|y) = \frac{p(x|z, y) p(z|y)}{p(z|x, y)} \quad (11)$$

could be evaluated directly, but $p(z|x, y)$ is also intractable.

A (conditional on y) variational autoencoder [7,8] is a deep-learning tool for treating such a latent variable model. It introduces a *recognition* (or *encoder*) model $q(z|x, y)$, which is an approximation to the posterior $p(z|x, y)$. As with the first two networks, the recognition network should have tractable density and be easy to sample, e.g., a multivariate Gaussian. One can then take the expectation (over z) of the logarithm of the posterior,

$$\begin{aligned} \log p(x|y) &= \mathbb{E}_{q(z|x, y)} \log p(x|y) \\ &= \mathbb{E}_{q(z|x, y)} \log \frac{p(x, z|y)}{q(z|x, y)} \\ &\quad + D_{\text{KL}}(q(z|x, y) \| p(z|x, y)) \\ &\equiv \mathcal{L} + D_{\text{KL}}(q(z|x, y) \| p(z|x, y)), \end{aligned} \quad (12)$$

where the last term is the KL divergence,

$$D_{\text{KL}}(q(z|x, y) \| p(z|x, y)) \equiv \mathbb{E}_{q(z|x, y)} \log \frac{q(z|x, y)}{p(z|x, y)}. \quad (13)$$

Since this is non-negative, \mathcal{L} is known as the *variational lower bound* on $\log p(x|y)$. If $q(z|x, y)$ is identical to $p(z|x, y)$, then $\mathcal{L} = \log p(x|y)$.

The variational autoencoder maximizes the expectation of \mathcal{L} over the true distribution. The associated loss function can be written

$$\begin{aligned} L_{\text{CVAE}} &= \mathbb{E}_{p_{\text{true}}(x)} \mathbb{E}_{p_{\text{true}}(y|x)} [-\mathcal{L}] \\ &= \mathbb{E}_{p_{\text{true}}(x)} \mathbb{E}_{p_{\text{true}}(y|x)} [-\mathbb{E}_{q(z|x, y)} \log p(x|z, y) \\ &\quad + D_{\text{KL}}(q(z|x, y) \| p(z|x, y))]. \end{aligned} \quad (14)$$

The three networks, $p(z|y)$, $p(x|z, y)$, and $q(z|x, y)$, are trained simultaneously. To evaluate the loss, the outer two expectation values are treated the same as in the previous subsection. The inner expectation value is evaluated using a Monte Carlo approximation, typically with a single sample from $q(z|x, y)$. For Gaussian $q(z|x, y)$ and $p(z|y)$ the KL divergence term may be calculated analytically; otherwise, a single Monte Carlo sample suffices.

For training, it is necessary to take the gradient of the loss with respect to network parameters. The stochasticity of the Monte Carlo integral estimates must, therefore, be carefully treated. This can be done by using the *reparameterization trick* [7,8], namely, by treating the random variable as an additional input to the network drawn from a fixed distribution. For example, if $q(z|x, y)$ is a multivariate Gaussian with mean $\mu(x, y)$ and Cholesky matrix $A(x, y)$, then

$$z = \mu(x, y) + A(x, y)\epsilon, \quad \text{with } \epsilon \sim \mathcal{N}(0, 1)^l, \quad (15)$$

is a sample from $q(z|x, y)$. With this trick, one can now take the gradient of z with respect to network parameters.

This setup is called a variational *autoencoder* because the first term in the loss function has the form of an autoencoder. The recognition network $q(z|x, y)$ is known as the encoder, and $p(x|z, y)$ as the decoder. This first term (the reconstruction loss) is minimized if x is recovered after being encoded into z and then decoded. The other term in the loss function (the KL loss) pushes the encoder to match the prior $p(z|y)$ and acts as a regulator. When the variational autoencoder works as an autoencoder, the latent variables z can give a useful low-dimensional representation of x .

Recent work [5] used a CVAE with diagonal Gaussian networks to model gravitational-wave posterior distributions, achieving excellent results over the four parameters (m_1, m_2, t_c, d_L) . In the following two subsections we describe the use of masked autoregressive flows to build even more general distributions.

C. Masked autoregressive flows

In this subsection we review the concept of a masked autoregressive flow, a type of normalizing flow that we use in our work to map simple distributions into more complex ones. We refer the reader to [10] for a much more in-depth discussion.

Consider a probability density $p(x)$. Without any loss of generality, this may be written using the chain rule as

$$p(x) = \prod_{i=1}^n p(x_i | x_{1:i-1}). \quad (16)$$

An autoregressive model restricts each conditional distribution in the product to have a particular parametrized form. We will take this to be univariate normal [10],

$$p(x_i | x_{1:i-1}) = \mathcal{N}(\mu_i(x_{1:i-1}), \exp(2\alpha_i(x_{1:i-1}))), \quad (17)$$

for $i = 1, \dots, n$. (The mean μ_i and variance $e^{2\alpha_i}$ that define this distribution are functions that will ultimately be given as outputs of a neural network.)

In [11] it was observed that an autoregressive model defines a normalizing flow. In other words, suppose $u \sim \mathcal{N}(0, 1)^n$. Then

$$x_i = \mu_i(x_{1:i-1}) + u_i \exp \alpha_i(x_{1:i-1}) \quad (18)$$

gives a sample from $p(x)$. This mapping $f: u \mapsto x$ is defined recursively in (18), but the inverse mapping,

$$u_i = [x_i - \mu_i(x_{1:i-1})] \exp(-\alpha_i(x_{1:i-1})), \quad (19)$$

is nonrecursive. Because f is autoregressive, the Jacobian determinant is very simple,

$$\left| \det \frac{\partial(f_1^{-1}, \dots, f_n^{-1})}{\partial(x_1, \dots, x_n)} \right| = \exp \left(- \sum_{i=1}^n \alpha_i(x_{1:i-1}) \right). \quad (20)$$

Hence, f defines a normalizing flow. Starting from a simple base distribution, the change of variables rule (2) can be used to evaluate the density $p(x)$.

An autoregressive flow can be embedded within a neural network [19]. For the affine flow (18), the network takes as input $x_{1:n}$, and gives as output $\mu_{1:n}$ and $\alpha_{1:n}$. To preserve the autoregressive property (i.e., that μ_i and α_i depend only on $x_{1:i-1}$), however, there can be no path from input x_i to output $\mu_{<i}$ or $\alpha_{<i}$. Such a network can be constructed by starting with a fully connected network (in general, including hidden layers) and applying a mask to the weight matrices in such a way that these paths are absent. An algorithm for constructing this mask is given in [19]; the resulting model is called a MADE (masked autoencoder density estimator).

For the inverse flow (19), a single pass through the MADE network is needed to obtain $\mu_{1:n}$ and $\alpha_{1:n}$, since all components $x_{1:n}$ are known [11]. The forward direction (18), however, requires n passes: to calculate μ_i and α_i , one requires $x_{1:i-1}$, which in turn requires $\mu_{1:i-1}$ and $\alpha_{1:i-1}$. Thus, the forward flow takes n times longer to calculate. So, although the flow is invertible, the model should be designed to use only the inverse flow during training, when many applications of the flow are needed. At evaluation time, however, the forward pass may be required.

For an autoregressive model such as MADE, the order of the components of x is important: components have distributions conditional on only lower-index components. Moreover, x_1 is independent of all other components, and follows a fixed normal distribution. To achieve sufficient generality, it is necessary to stack several MADEs in sequence, permuting the order of the components between each pair [11]. This is called a MAF [10]. For stability during training it is also useful to insert batch normalization layers [20] between the MADE blocks, and between the base distribution and the first MADE block; in the context of a flow, these also contribute to the Jacobian [10]. MAFs and related networks have been used to model very complex distributions over high-dimensional spaces, including audio [21] and images [22].

In the context of gravitational-wave parameter estimation, the sample space is relatively low dimensional, $n = \mathcal{O}(10)$. To model the posterior distribution, each MADE block must be made conditional on the (high-dimensional) strain data y , while maintaining the autoregressive property over x . In other words, the masked networks that define the MADEs take y as additional input, but no masking is applied between y and the first hidden layer of the network. We can then take a standard normal base distribution, and flow it through all the MADE blocks and batch normalization layers, to obtain the complex posterior. The loss function is the same as in

Sec. II A, but with the change of variables rule used to evaluate the density, i.e.,

$$L = \mathbb{E}_{p_{\text{true}}(x)} \mathbb{E}_{p_{\text{true}}(y|x)} \left[-\log \mathcal{N}(0, 1)^n(f_y^{-1}(x)) - \log \left| \det \frac{\partial(f_{y,1}^{-1}, \dots, f_{y,n}^{-1})}{\partial(x_1, \dots, x_n)} \right| \right], \quad (21)$$

where now f_y denotes the entire sequence of flows in the network, and we have made explicit the dependence on strain data y . Notice that this involves only the inverse flow f_y^{-1} , which, compared to the forward flow f_y , is fast to evaluate because it does not involve recursion.

D. Combined models

More powerful models can be obtained by combining autoregressive flows with the variational autoencoder. Each of the three networks comprising the CVAE—the encoder $q(z|x, y)$, the decoder $p(x|z, y)$, and the prior $p(z|y)$ —can be made more flexible by applying autoregressive flows to their sample spaces. We discuss each of these possibilities in turn. In our experiments, we found that the best performance was achieved when combining all three. In all cases, the CVAE loss function (14) is optimized, with the change of variables rule used to evaluate the component densities.

1. Encoder with inverse autoregressive flow

Normalizing and autoregressive flows were first proposed as a way to increase the flexibility of the encoder network [9, 11]. This is important because the CVAE loss function (14) differs from the desired cross-entropy loss (8) by the expectation of the KL divergence between $q(z|x, y)$ and the intractable $p(z|x, y)$. If this can be made smaller, then the two losses converge; hence, $q(z|x, y)$ should be as flexible as possible.

A flexible encoder is also desired to avoid a situation called “posterior collapse.” The reconstruction and KL loss terms are in competition during training, and if the KL loss term collapses to zero, the network can fail to autoencode. In this situation, the encoder matches the prior, so it ignores its x input; the latent variables z contain no information about x beyond what is contained in y . This can happen either because the loss gets stuck in an undesired local minimum during training, or the configuration with vanishing KL loss is actually the global minimum [12]. The former can be alleviated by careful training strategies such as annealing the KL loss term [23]. The latter can occur because the use of latent variables incurs a cost related to $D_{\text{KL}}(q(z|x, y) || p(z|x, y))$ [12]. If the decoder is sufficiently powerful such that it can model $p(x|y)$ on its own, then $p(z|x, y) \rightarrow p(z|y)$ and $L_{\text{CVAE}} \rightarrow L$ is the global minimum. The network simply decides it is not worthwhile to use the latent variables.

Thus, a flexible encoder is important for performance and to make full use of latent variables. To make the

encoder distribution more flexible, one could start with a Gaussian encoder, as in Sec. II B, and then apply a normalizing flow to the sample space. Since evaluation of L_{CVAE} requires sampling from $q(z|x, y)$, an *inverse autoregressive flow* (IAF) should be used [11]. Sampling is fast since the inverse flow (19) does not involve recursion. The density evaluation needed for the KL loss term only needs to take place for z sampled from $q(z|x, y)$, so caching may be used.

Each MADE block comprising the encoder IAF may be conditioned on x and y . In addition, one is free to output an additional context variable w from the network that specifies the normal distribution, and condition the IAF on this. (The reason for this additional variable is simply to provide additional information paths from the encoder inputs x and y to the IAF, in order to improve training. The initial work [11] conditioned the IAF only on w .) In our experiments we found it most effective to condition only on w and x .

2. Decoder with masked autoregressive flow

To increase the flexibility of the Gaussian decoder of Sec. II B, one can similarly apply a normalizing flow to its sample space. However, this time a forward MAF should be used, because at train time, a density evaluation of $p(x|z, y)$ is required to evaluate the reconstruction loss term in L_{CVAE} . The MADE blocks may be conditioned on y (as in the basic MAF of Sec. II C) and the latent variable z . This powerful decoder increases the risk of posterior collapse, so it is useful to use also the IAF encoder.

3. Prior with masked autoregressive flow

As shown in [12], a normalizing flow applied to the sample space of the prior network $p(z|y)$ effectively adds flexibility to both the encoder network and the decoder network. For fast training, one should again use a forward MAF because a density evaluation of $p(z|y)$ is needed to compute the KL loss term in L_{CVAE} . The MADE blocks should be conditioned on y .

Although a prior MAF and an encoder IAF are very closely related, in that both transform the latent space, they differ in that the encoder IAF can be conditioned on x . We found that this had a significant impact on performance, and therefore included both autoregressive flows in our models.

III. NONSPINNING BINARIES

In this section we describe our experiments in using deep-learning models to describe the posteriors of non-spinning coalescing black hole binaries. Binaries are described by five parameters: the masses m_1 and m_2 , the luminosity distance d_L , the time of coalescence t_c , and the phase of coalescence ϕ_0 . These parameters and their ranges are chosen to facilitate comparison with [5].

A. Training data

Training data for our models consist of pairs of parameters x and strain time series y . Parameters are sampled from a prior distribution $p(x)$, which is uniform over each parameter except for the volumetric prior over d_L , $p(d_L) \propto dV_c/dd_L$, where V_c is the comoving volume. This volumetric prior assumes sources are distributed uniformly in comoving volume and is the standard prior used in LIGO/Virgo analyses. In practice there is likely to be evolution in the rate of mergers which will mean the true distribution in luminosity distance is different, but it is straightforward to reweight posterior samples to alternative priors as required. Parameter ranges are taken from [5],

$$35 M_\odot \leq m_1, m_2 \leq 80 M_\odot, \quad (22a)$$

$$1000 \text{ Mpc} \leq d_L \leq 3000 \text{ Mpc}, \quad (22b)$$

$$0.65 \text{ s} \leq t_c \leq 0.85 \text{ s}, \quad (22c)$$

$$0 \leq \phi_0 \leq 2\pi. \quad (22d)$$

We also take $m_1 \geq m_2$.

Strain realizations are in the time domain and 1 s long ($0 \leq t \leq 1$ s) with a sampling rate of 1024 Hz. (We found that the sampling rate of 256 Hz of [5] was not sufficient to eliminate Gibbs ringing.) Strain data consist of a waveform $h(x)$, deterministic from parameters x , and random noise n , sampled from the Advanced LIGO [24] Zero Detuned High Power PSD [25]. In this analysis we consider the response of a single gravitational-wave detector only. With a single detector it is not possible to determine the extrinsic parameters of the gravitational-wave source that describe the location of the source on the sky, its polarization, and its orientation with respect to the detector. There is therefore no point in including these parameters in the waveform model. Therefore, we take $h(x)$ to be the “+” polarization of IMRPhenomPv2 waveforms [26], which, following [5], we whitened in the frequency domain using the PSD before taking the inverse Fourier transform to the time domain. Finally, we rescaled our waveforms by dividing by $\sqrt{2\Delta t}$ so that in the end, time domain noise is described by a standard normal distribution in each time sample, i.e., $n \sim \mathcal{N}(0, 1)^m$.

The whitening and rescaling procedure serves two purposes: it means that we can easily draw noise realizations at train time by sampling from a standard normal distribution, and it ensures that the input data to the neural network has approximately zero mean and unit variance. Similarly, we rescale x to have zero mean and unit variance in each component across the training set. This “standardization” of training data is a standard approach, which improves training speed and ensures all inputs are equally weighted initially [27].




FIG. 1. Sample waveform h and strain realization y . The SNR for this injection is 24.9.

Our dataset consists of 10^6 (x, h) pairs, which we split into 90% training data and 10% validation data. Noise realizations n are sampled at train time to give strain data $y = h + n$; a sample time series is given in Fig. 1. The median signal-to-noise ratio (SNR) of our training set is 25.8, and the complete SNR distribution is given in Fig. 2. Our dataset is the same size as that of [5] and a factor of 10^3 smaller than the effective dataset of [4]. Larger datasets are generally preferred to prevent overfitting, but they are also more costly to prepare and store. We would like to build a system that can generalize well with as small a dataset as possible in order to minimize these costs when moving to more complicated and longer waveforms in the future. By keeping track of performance on the validation set when training our models, we found that our training set was sufficiently large to avoid overfitting. We also experimented with a 10^5 -element dataset; this had a slightly reduced, but still acceptable, performance.

B. Experiments

We modeled our gravitational-wave posteriors with the three types of neural network described in Sec. II: a CVAE (similar to [5]), a MAF, and a CVAE with autoregressive




FIG. 2. Histogram of SNR values for our training set.

TABLE II. Network depth, number of trainable parameters, and training time. Numbers of trainable parameters are approximated as the numbers of weights; for autoregressive flows, we included an estimate of the number of *unmasked* weights.

Model	Layers	Parameters ($\times 10^6$)	Train time (h)
CVAE	12	31.6	5.6
MAF	20	4.0	6.5
CVAE+	72	18.8	14.8

TABLE III. Final values of training and validation set loss functions after 250 epochs. Lower overall values of the loss indicate better performance, and small differences between training and validation loss indicate that the networks are not overfitting the training data. Loss values are not directly comparable to [5] due to differences in data standardization.

Model	Train loss	Val loss	KL train loss	KL val loss
CVAE	-4.52	-4.31	4.33	4.49
MAF	-4.43	-4.42
CVAE+	-7.00	-6.95	6.76	6.77

flows appended to the three subnetworks (denoted CVAE+). We selected hyperparameters based on choices in the literature [5,10] and on sweeps through hyperparameter space. Approximate network sizes and training times for each architecture are listed in Table II, and final loss functions after training in Table III.

All of our networks use rectified linear unit nonlinearities. They were trained for 250 epochs with a batch size of 512, using the Adam optimizer [28]. The learning rate was reduced by a factor of 2 every 80 epochs. We trained all networks for the same number of epochs in order to have a fair comparison, and we felt that 250 epochs represented a reasonable tradeoff between performance gain and training time. After 250 epochs the learning rate has been reduced by a factor of 8, significant improvements in performance are unlikely and overfitting becomes increasingly probable. Sampling performance results for the three network architectures are collected in Figs. 3 and 4.

I. CVAE

Our CVAE network is designed to be similar to that of [5]. The encoder, decoder, and prior distributions are all taken to be diagonal Gaussians, each parametrized by fully connected neural networks with three hidden layers of dimension 2048. The latent space is of dimension $l = 8$. We used the same initial training rate as [5], of 0.0001.

To train the CVAE, we optimize the variational lower bound loss L_{CVAE} of (14), with a single Monte Carlo sample to estimate the expectation value over $q(z|x, y)$. Since L_{CVAE} is only an upper bound on the cross-entropy loss L , the value of the loss function alone is not entirely indicative of performance. Indeed, when posterior collapse

occurs (see Sec. II D 1), the gap between L_{CVAE} and L can vanish; in this case, the value of L is larger than that of a network with the same L_{CVAE} and no posterior collapse. For this reason, it is important to also use other metrics to evaluate performance. For CVAE models we always quote the KL loss as an indication that the latent space is being used.

To encourage use of the latent space, we found it to be beneficial to use KL annealing during training, i.e., we multiplied the KL loss part of L_{CVAE} by a factor between zero and one during the early stages. This reduces the importance of the KL loss term compared to the reconstruction loss. For all CVAE-based models, we adopted a cyclic KL annealing strategy [29]: for the first 6 epochs we used annealing factors of $(10^{-5}, 1/3, 2/3, 1, 1, 1)$, and we repeated this cycle 3 more times; see Fig. 5. We also ignored the KL loss term whenever it was less than 0.2.

In Fig. 3(a), we show a posterior distribution corresponding to the strain data y of Fig. 1. This is constructed from $N = 5 \times 10^4$ posterior samples, obtained using formula (10) with the prior and decoder networks as follows: (1) pass y through the prior network to obtain the distribution $p(z|y)$, (2) draw N latent-space samples $z^{(i)}$ from the prior, (3) pass these through the decoder network to obtain N distributions $p(x|z^{(i)}, y)$, and finally (4) draw one sample $x^{(i)}$ from each of these distributions.

By inspection of the posterior, it is clear that the latent space is being used in the model, since the distribution is not a diagonal Gaussian. The distributions for most of the parameters look reasonable, and they cover the true values of the parameters. The phase of coalescence, ϕ_0 , is, however, not being captured at all. Indeed, ϕ_0 should be precisely π periodic because our training set waveforms only contain the $(l, m) = (2, 2)$ mode of the signal. Moreover, since we are taking a single polarization, ϕ_0 should be resolvable.

We can evaluate the performance of the network with a p-p plot [30]. To do this, we compute posterior distributions, each comprised of 10^4 samples, for 10^3 different strain realizations [i.e., $y = h(x) + n$ for x drawn from the prior over parameters and n a noise realization]. For each one-dimensional posterior, we then compute the percentile within the posterior at which the true parameter lies. For each parameter, the p-p curve is then the cumulative density function of the percentile values. This is shown in Fig. 4(a) for the CVAE network. If the CDF is diagonal, then the percentile values are distributed according to a uniform distribution, as one would expect for the true one-dimensional posteriors. We can see that the CVAE appears to capture all of the one-dimensional distributions except for ϕ_0 .

To confirm that the percentile scores are well distributed, we also performed a Kolmogorov-Smirnov test. We calculated the KS statistic to compare the distribution of percentile scores to a uniform distribution. We found that ϕ_0 had miniscule p value, as expected, the p value of t_c was

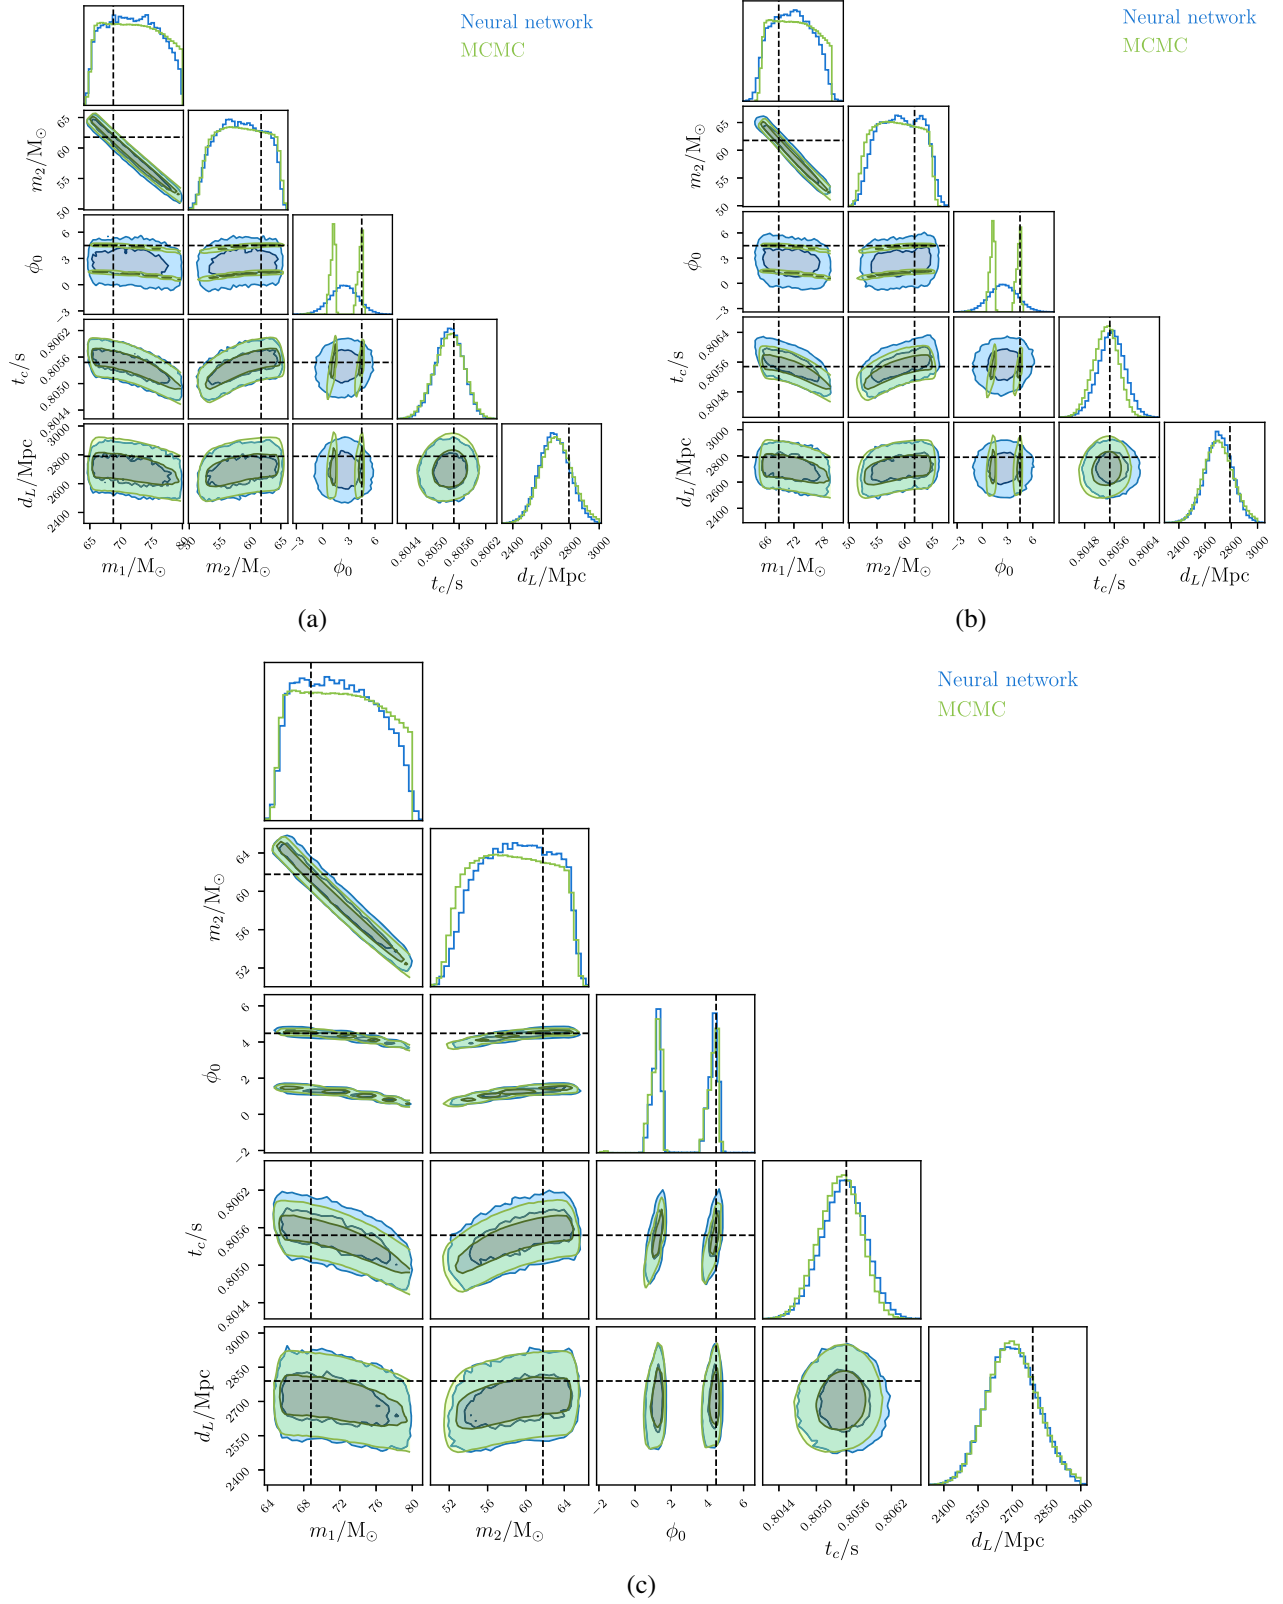


FIG. 3. One- and two-dimensional marginalized posterior distributions for the strain realization shown in Fig. 1, comparing output from (a) CVAE, (b) MAF, and (c) CVAE+ neural networks. Each figure is constructed from 5×10^4 neural network samples, and contours represent 50% and 90% credible regions. MCMC results are overlayed for comparison. The CVAE+ network is the only one capable of capturing the multimodality in the ϕ_0 posterior.

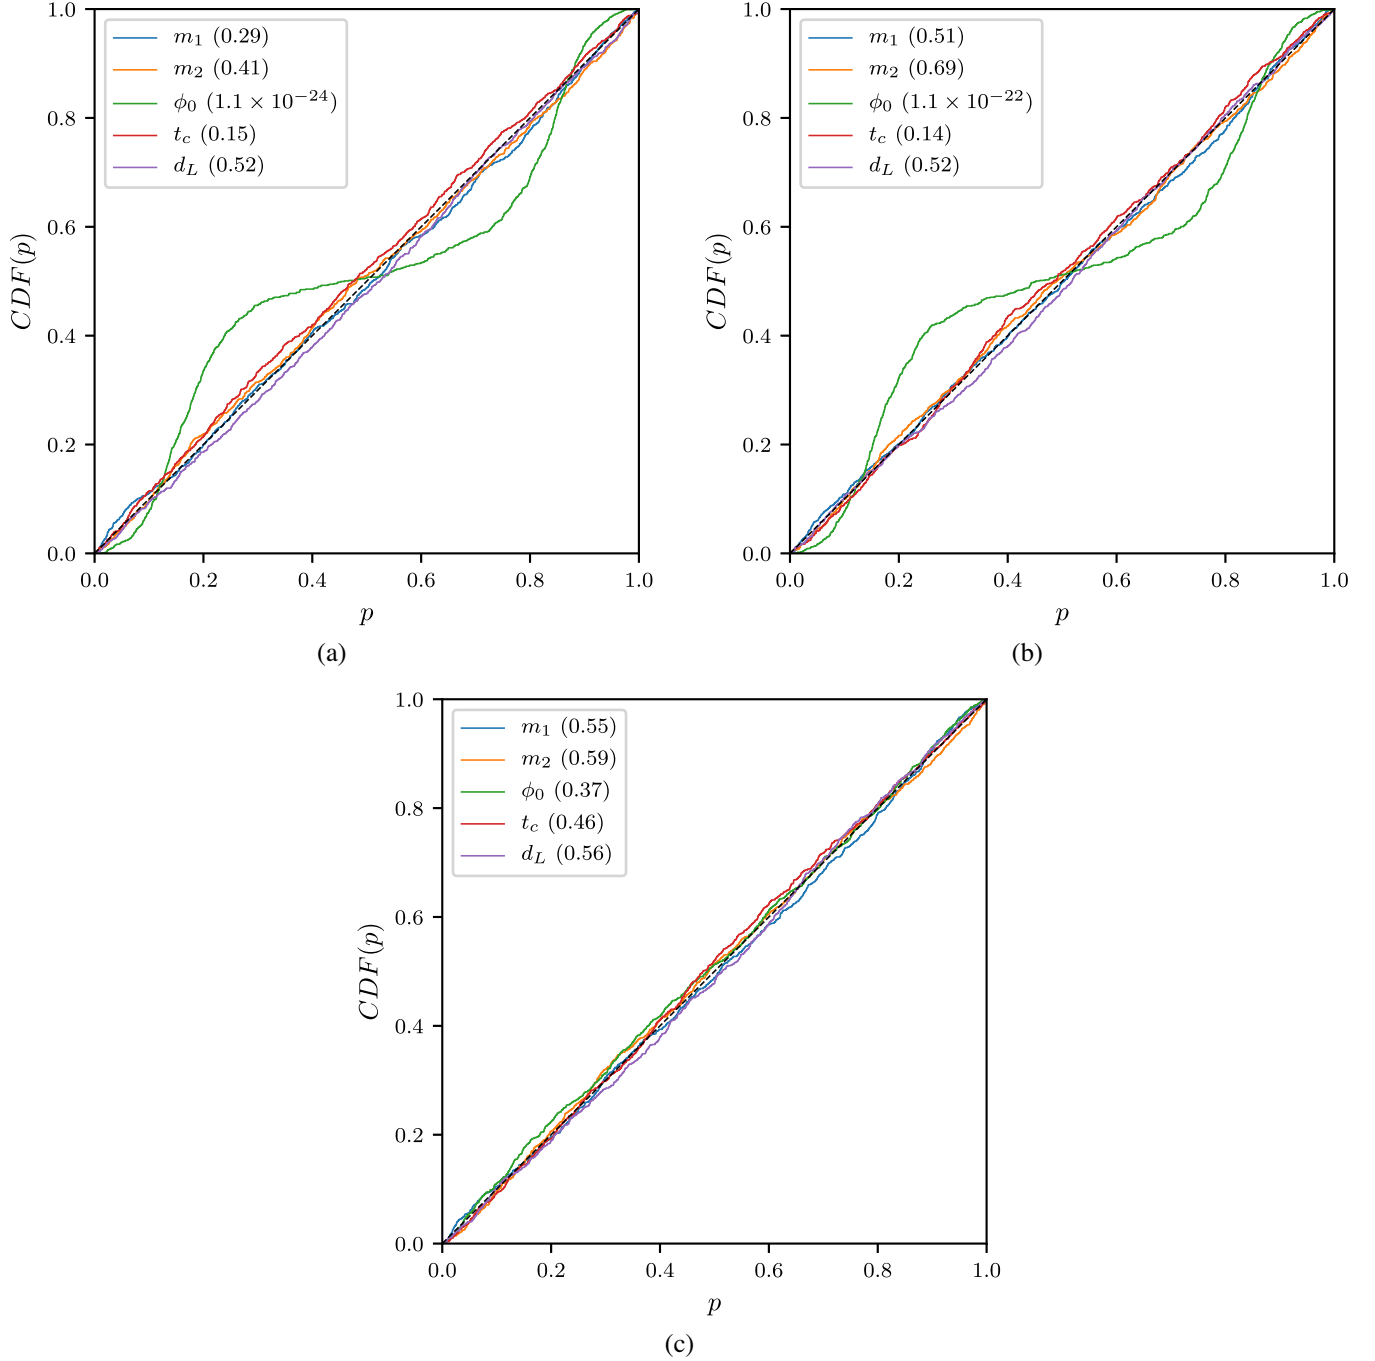


FIG. 4. p - p plots for (a) CVAE, (b) MAF, and (c) CVAE+ neural network models. For each marginalized one-dimensional posterior distribution, the curve represents the cumulative distribution function of the percentile scores of the true values of the parameter. Each plot is constructed from the same 10^3 parameter choices and strain realizations. Deviations from the diagonal represent a failure of the model to learn the true posterior for that parameter. KS test p values are provided in the legends.

0.15, and all other p values were greater than 0.29. This indicates that all parameters are well recovered, except for ϕ_0 .

2. MAF

Next, we modeled the gravitational-wave posterior $p(x|y)$ directly using a masked autoregressive flow. As

described in Sec. II C, the MAF network describes the posterior in terms of an invertible mapping f_y from a five-dimensional standard normal distribution $\mathcal{N}(0, 1)^5(u)$ into the gravitational-wave posterior $p(x|y)$. The flow $f_y(u)$ is autoregressive over u and has arbitrary dependence on strain data y . The MAF network does not involve latent variables, and optimizes the cross entropy with the data




FIG. 5. Training and validation loss for each epoch for the CVAE+ network. Spikes at early stages arise from the cyclic KL annealing. Since training and validation loss coincide, there is negligible overfitting. The qualitative behavior of the CVAE and MAF networks is similar, so we reported only final values in Table III.

distribution. Thus the loss function L [given in (21)] can be used to directly compare performance for different models.

Our best performing MAF network consists of five MADE blocks, each with three hidden layers of dimension 512, and conditioned on y . We also inserted batch normalization layers between each pair of MADE blocks, and between the first MADE block and the base space. We found this to be important for training stability. We were able to train the MAF successfully at a higher initial learning rate than the CVAE, of 0.0004. Following [10], at the end of each training epoch, before computing the validation loss, we set the stored mean and variance vectors of the batch normalization layers to the mean and variance over the entire training set. (We did this also for the CVAE+ network below.)

In Fig. 3(b), we show a corner plot for the same strain data as before. Samples are obtained by first sampling from the standard normal base space, $u \sim \mathcal{N}(0, 1)^5$ and then applying the flow $f_y(u)$. All quantities appear to be well modeled except, again, for ϕ_0 .² This is consistent with the p-p plot shown in Fig. 4(b). The KS statistic p values were 0.14 for t_c , and they were greater than 0.5 for all other parameters (except for ϕ_0). (The p-p plots for all networks are computed from the same strain realizations, so it is consistent to see the same KS statistic p values for the different networks.)

The performance of the MAF network is comparable to that of the CVAE described in [5], but with one eighth the

²In some experiments with deeper MAF networks, we were in fact able to properly resolve the ϕ_0 posterior. This was, however, difficult to reproduce and somewhat unstable because MAF networks use element-wise affine transformations (18), and these may not be flexible enough to consistently resolve multimodality. We leave further investigation to future work.

number of trainable parameters.³ In both cases, all parameters except for ϕ_0 are well modeled by the networks. In addition, the final loss values given in Table III are very close for both networks, with a slight edge in validation loss for the MAF. However, since L_{CVAE} is an upper bound on L , the cross-entropy loss for the CVAE network may actually be lower than that of the MAF. Indeed, the fact that the CVAE made use of the latent space suggests that this gap is nonzero. Comparison of the training and validation loss functions shows slight overfitting for the CVAE, but none for the MAF.

3. CVAE+

Finally, we experimented with combinations of CVAE and MAF networks. As described in Sec. IID, all three component distributions of the CVAE can be made more flexible by applying MAF transformations to the diagonal Gaussian distributions, thereby increasing the total modeling power of the network [11,12]. Indeed, Appendix A of [11] shows that a single linear autoregressive flow is capable of generalizing a Gaussian distribution from diagonal to generic covariance.

For the combined models, the initial Gaussian distributions (the base spaces for the autoregressive flows) are modeled in the same way as in Sec. IIIB 1, as fully connected three-hidden-layer networks. However, we reduced the size of hidden layers to 1024 dimensions. We also kept the $l = 8$ dimensional latent space. We found that best performance was achieved by applying autoregressive flows to all three component distributions as follows:

$q(z|x, y)$: We added an IAF after the initial Gaussian encoder. This was made conditional on x and an extra eight-dimensional context output w from the initial encoder; see Sec. IID 1. We also experimented with conditioning on y , but found this to reduce performance.

$p(x|y)$: We added a MAF after the initial Gaussian prior network. This was made conditional on y .

$p(x|z, y)$: We added a MAF after the initial Gaussian decoder network. This was conditional on y and z .

In all cases the MAF/IAF parts consist of five MADE blocks (three hidden layers of 512 dimensions each) and batch normalization layers. Although the CVAE+ network contains far more layers than the basic CVAE, it has roughly half the number of free parameters since we reduced the width of the hidden layers in the Gaussian components.

Training was performed by optimizing L_{CVAE} , with the change of variables rule used to evaluate the component

³It is possible that the number of trainable parameters in the CVAE could be reduced without a significant decrease in performance, but we assume that [5] optimized hyperparameters appropriately.

densities. Sampling was similar to the basic CVAE, but now with the appropriate flows applied after sampling from the Gaussians.

We found that optimization was best when we trained the Gaussian distributions at a learning rate of 0.0001, and the autoregressive networks at a higher rate of 0.0004, combining the rates of the previous two subsections. As always, the learning rate was reduced by half every 80 epochs.

It is especially important for the CVAE+ to apply some sort of KL annealing to encourage use of the latent space. An important difference compared to the basic CVAE of Sec. III B 1 is that now the decoder network is sufficiently powerful to model much of the posterior *on its own*. Indeed, it is just as flexible as the MAF network of the previous subsection, which produced the posterior in Fig. 3(b), and it is well known that a CVAE will often ignore latent variables if the decoder is sufficiently powerful [12,23]. This strategy combined with the flexible encoder distribution resulted in higher KL validation loss for the CVAE+ network (6.77) than the basic CVAE (4.49) by the end of training; see Table III. A plot showing the KL and total loss terms as training proceeds is given in Fig. 5.

A sample gravitational-wave posterior distribution for CVAE+ is given in Fig. 3(c). In contrast to the simpler models, this captures the periodicity in ϕ_0 . Sampling was very fast, requiring ≈ 0.7 s to obtain 10^4 samples. The p-p plot in Fig. 4(c) shows excellent statistical performance, with all KS statistic p values greater than 0.35.

For validation of our network against standard methods, Fig. 3 also includes samples obtained using MCMC. We used the emcee ensemble sampler [15] with the same prior and likelihood as defined by our training set construction. The ensemble sampler stretch move scale parameter was chosen to be $a = 1.4$, we used 500 walkers for 70,000 steps, and we finally thinned samples by a factor of 10. There is clear agreement between MCMC sampling and neural network sampling for the CVAE+ network in Fig. 3(c), with slight deviation in the mass posteriors. This might be partially due to sampling fluctuations in the MCMC results, but we expect that any residual differences in the CVAE+ results could be reduced with hyperparameter improvements or further training.

The KL divergence between the two distributions shown in Fig. 3(c) was estimated using the method described in [31]. Because of the computational cost of the approach, we computed the KL divergence using 10,000 samples drawn from each posterior. The exact value obtained depends on the particular samples drawn, but it was typically of order 0.3. To put this in context, we used the same approach to compute KL divergences between two sets of samples drawn from the same posterior (either the neural network posterior or the MCMC posterior). These KL divergences were less than but of the order of 0.1. This is consistent with the slight deviations between the neural network and MCMC posteriors seen in Fig. 3(c).

IV. INCLINED BINARIES WITH ALIGNED SPINS

To test our method with a more challenging dataset, we augmented the parameter space to include nonzero aligned spins and inclination angle. We took uniform priors on χ_{1z}, χ_{2z} , and $\cos \theta_{JN}$ between -1 and 1 , keeping the rest of the prior distribution unchanged. The neural network was trained to model posterior distributions over the eight-dimensional parameter space $x = (m_1, m_2, \phi_0, t_c, d_L, \chi_{1z}, \chi_{2z}, \theta_{JN})$. We held the dataset size fixed at 10^6 elements, and again split 90% into the training set and 10% into the validation set. The median SNR for this training set is 17.2, the mean is 19.1, and the range is (6.1, 100.9).

We tested only the CVAE+ model, since this performed best in previous experiments. Because gravitational-wave posteriors over the eight-dimensional parameter space have increased complexity, we increased the capacity of our network by doubling the latent-space dimension to $l = 16$, the number of MADE blocks to 10 per component network, and the dimension of the IAF context variable to 16. We also found performance was best when we froze the Gaussian part of the prior network, but aside from this all other hyperparameters were unchanged. This increased the total number of trainable parameters to 25.2×10^6 .

We trained again for 250 epochs, with final total loss values of -0.83 (train) and -0.73 (validation), and final KL loss values of 11.05 (train) and 11.14 (validation), showing very little overfitting. This also indicates that for the extended parameter space, the network made heavier use of the latent space. (This observation motivated the increased size of the latent space.)

In Fig. 6 we show a representative posterior distribution produced by the neural network, as well as a corresponding MCMC posterior. The neural network continues to capture the m_1 - m_2 degeneracy, as well as the new χ_{1z} - χ_{2z} and d_L - θ_{JN} degeneracies that arise in the extended parameter space. The distribution over masses, however, is slightly too broad in the neural network posterior. With improvements to network architecture, or additional training, we expect further convergence of the neural network to MCMC. We note also that in contrast to the smaller parameter space of the previous section, the posterior over ϕ_0 is simply the prior. The reason for this is that in our waveform model, there is a three-parameter degeneracy between the two spins χ_{1z}, χ_{2z} , and the phase ϕ_0 , where small changes in the spins cause large changes in the phase.

The inset in Fig. 6 shows the posterior distribution over the chirp mass $\mathcal{M} = (m_1 m_2)^{3/5} / (m_1 + m_2)^{1/5}$, asymmetric mass ratio $q = m_2 / m_1$, and effective spin parameter $\chi = (m_1 \chi_{1z} + m_2 \chi_{2z}) / (m_1 + m_2)$. These are derived quantities, with sampling done instead over the parameters listed above. Although posteriors are simpler over the derived parameters, to test our method we chose to sample over parameters with more nontrivial posteriors that include non-Gaussianities and correlations between parameters.

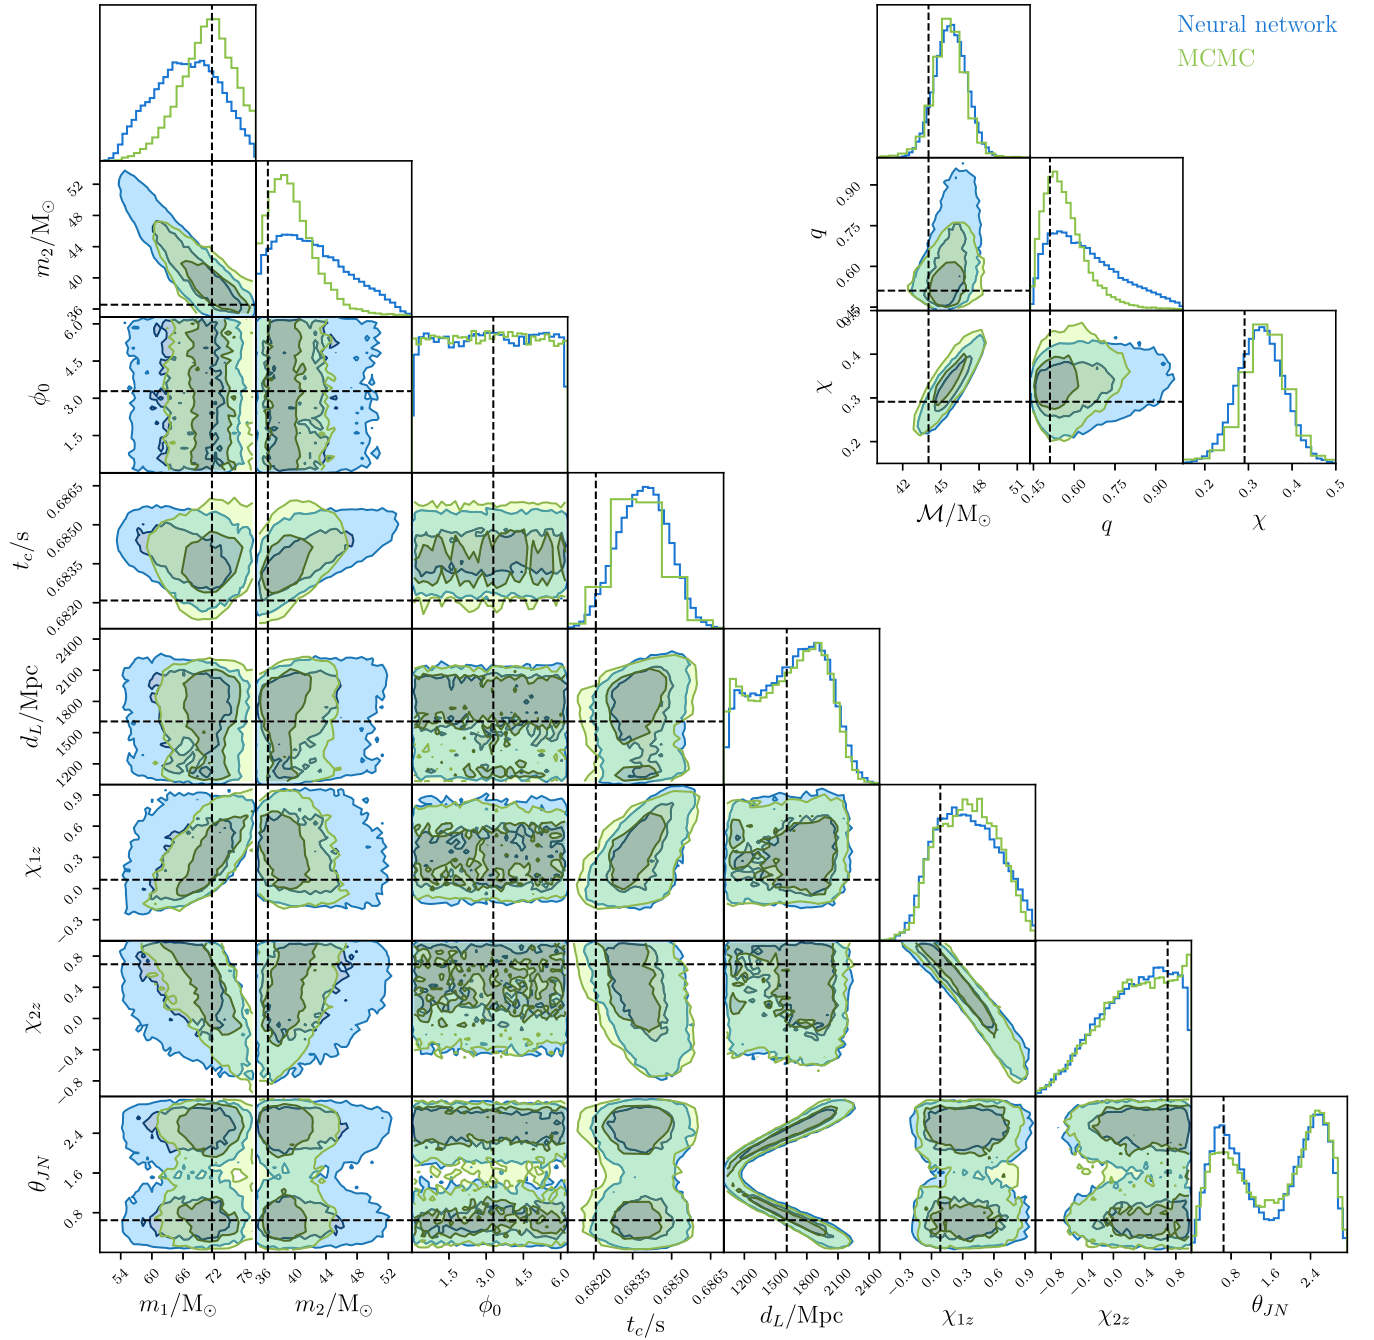


FIG. 6. Sample posterior distribution for eight-dimensional parameter space, from 5×10^4 samples of CVAE+. Inset shows derived quantities: chirp mass \mathcal{M} , asymmetric mass ratio q , and effective spin parameter χ . The waveform injection has a SNR of 29.6.

We see that in comparison to MCMC, \mathcal{M} and χ are well recovered, but the neural network posterior over q is slightly too broad, reflecting the broad m_1 – m_2 posterior in the main figure. A p–p plot for the extended parameter space is presented in Fig. 7. By inspection, this shows good recovery of all parameters. Moreover, all parameters have KS p values greater than 0.2. Although in the particular case of Fig. 6, the mass-ratio posterior does not agree as well with the emcee results as the other parameters, the

posterior clearly differs from the prior, and the p–p plots for the masses are diagonal. Therefore these differences are not affecting the overall performance of the network, but they are nevertheless an area in which improvements could be made in the future.

Sampling from the larger CVAE+ model used for the extended parameter space is slightly slower than the smaller model of the previous section, now requiring ≈ 1.6 s to obtain 10^4 posterior samples. This is partially due to the




FIG. 7. p - p plot for one-dimensional posteriors for eight-dimensional parameter space, modeled with the CVAE+ network. Constructed from 10^3 injections.

larger 16-dimensional latent space: the forward pass through a MAF is recursive, so twice as many passes are required to sample from the variational prior and decoder. Both MAFs also have twice as many layers.

V. CONCLUSIONS

In this work we introduced the use of masked autoregressive flows for improving deep learning of gravitational-wave posterior distributions. These learnable mappings on sample spaces induce transformations on probability distributions, and we used them to increase the flexibility of neural network posterior models. The models that we built can all be rapidly sampled, requiring <2 s to obtain 10^4 samples.

For nonspinning, quasicircular binary black holes, and a single gravitational-wave detector (a five-dimensional parameter space) we compared models involving a single MAF, a CVAE, and a CVAE with autoregressive flows applied to the encoder, decoder, and prior networks (CVAE+). We found that the performance of the single MAF and CVAE models were comparable, and that best performance was achieved by the CVAE+ model. The CVAE+ model was able to capture the bimodality in the phase ϕ_0 , which eluded the other models.

We then considered a larger eight-dimensional parameter space, with aligned spins and nonzero binary inclination angle. With a higher-capacity CVAE+ network, we successfully recovered the posterior distribution over all parameters, except for the mass ratio, which was slightly too broad. This demonstrates that our approach extends to

higher-dimensional parameter spaces. A modest increase in network capacity may, however, be required.

Although best performance was achieved with the CVAE+ model, an advantage of models without latent variables (such as the MAF alone) is that it is possible to directly evaluate the probability density function. Since the posterior distribution that the MAF models is normalized, one could then calculate the Bayesian evidence by separately evaluating the likelihood and prior. (In CVAE+, this would require marginalization over z .) Moreover, since the MAF loss function is just the cross-entropy loss with the true distribution, this means that loss functions of competing models can be compared directly. It would therefore be worthwhile to also try to improve the performance of the basic MAF model.

In contrast to typical applications of these deep-learning tools, the space Y on which all of our models are conditioned is of much higher dimensionality than the space X that we are modeling. One way to improve performance further may be to introduce convolutional neural networks to preprocess the strain data y and compress it to lower dimensionality. In the future, when we extend our models to treat longer waveforms, higher sampling rates, and additional detectors, this will become even more important.

Going forward, it will also be important to understand better the uncertainty associated to the neural network itself, particularly in regions of parameter space that are not well covered by training data. The training data is drawn from the prior and so, in principle, regions of the parameter space that have few samples in the training data are less likely to be observed astrophysically. However, the assigned priors are just reasonable guesses and the actual distribution of observed events may well differ from these assumptions. In addition, parts of the parameter space may have greater variation in the waveforms than others and therefore require a higher density in the training set to be properly learnt by the network. If the goal is to obtain unbiased posterior inference for every event observed then we might want to use training data drawn from a different prior that ensures equal accuracy over all possible parameters, and reweight the posterior samples to the desired prior at the end. In addition, or as an alternative, we need to obtain uncertainty estimates for the neural network output. Rather than taking the maximum likelihood estimate for the neural network parameters, as we did in this work, it is possible to model them as random variables with some probability distribution. This distribution can be learned through variational inference, and ultimately marginalized over, resulting in a slightly broader posterior over system parameters [6,32]. In our work, overfitting was not a problem, but such approaches to neural network uncertainty could be useful in situations where the binary system parameter space is not well covered due to high waveform generation costs.

As the rate of detected gravitational-wave events grows with improved detector sensitivity, methods for rapid parameter estimation will become increasingly desirable. For deep-learning approaches to become viable alternatives to standard inference methods, they must be extended to cover the full space of binary system parameters, and to treat longer duration waveforms from multiple detectors with detector PSDs that vary from event to event. Of particular interest is to extend this work to include inference of the sky location, as this will be beneficial for triggering multimessenger follow up of interesting events. There are no obvious conceptual obstacles to extending the network to inference of the extrinsic parameters. The dimensionality of the parameter space increases, as does the dimensionality of the input space, since such inference relies on using data from multiple detectors. As we saw in this paper, increasing

the dimensionality does increase the difficulty of designing and training the neural network, but not insurmountably and this is an obvious direction for further work. The methods discussed in this paper are an important first step toward the long term goal of building a rapid parameter estimation tool for characterizing arbitrary gravitational-wave detections.

ACKNOWLEDGMENTS

We would like to thank R. Cotesta, A. Ghosh, A. Matas, and S. Ossokine for helpful discussions. C. S. was supported by the EPSRC Centre for Doctoral Training in Data Science, funded by the U.K. Engineering and Physical Sciences Research Council (Grant No. EP/L016427/1) and the University of Edinburgh.

-
- [1] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs, *Phys. Rev. X* **9**, 031040 (2019).
 - [2] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Properties of the Binary Neutron Star Merger GW170817, *Phys. Rev. X* **9**, 011001 (2019).
 - [3] B. P. Abbott *et al.*, Multi-messenger observations of a binary neutron star merger, *Astrophys. J. Lett.* **848**, L12 (2017).
 - [4] A. J. K. Chua and M. Vallisneri, Learning Bayesian Posteriors with Neural Networks for Gravitational-Wave Inference, *Phys. Rev. Lett.* **124**, 041102 (2020).
 - [5] H. Gabbard, C. Messenger, I. S. Heng, F. Tonolini, and R. Murray-Smith, Bayesian parameter estimation using conditional variational autoencoders for gravitational-wave astronomy, [arXiv:1909.06296v2](https://arxiv.org/abs/1909.06296v2).
 - [6] H. Shen, E. A. Huerta, Z. Zhao, E. Jennings, and H. Sharma, Deterministic and Bayesian neural networks for low-latency gravitational wave parameter estimation of binary black hole mergers, [arXiv:1903.01998](https://arxiv.org/abs/1903.01998).
 - [7] D. P. Kingma and M. Welling, Auto-encoding variational Bayes, [arXiv:1312.6114](https://arxiv.org/abs/1312.6114).
 - [8] D. J. Rezende, S. Mohamed, and D. Wierstra, Stochastic backpropagation and approximate inference in deep generative models, in *International Conference on Machine Learning* (Proceedings of Machine Learning Research, Beijing, 2014), pp. 1278–1286, <http://proceedings.mlr.press/v32/rezende14.html>.
 - [9] D. Rezende and S. Mohamed, Variational inference with normalizing flows, in *International Conference on Machine Learning* (Proceedings of Machine Learning Research, Lille, 2015), pp. 1530–1538, <http://proceedings.mlr.press/v37/rezende15.html>.
 - [10] G. Papamakarios, T. Pavlakou, and I. Murray, Masked autoregressive flow for density estimation, in *Adv. Neural Inf. Process. Syst.* (2017), pp. 2338–2347, <http://papers.nips.cc/paper/6828-masked-autoregressive-flow-for-density-estimation>.
 - [11] D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling, Improved variational inference with inverse autoregressive flow, in *Adv. Neural Inf. Process. Syst.* (2016), pp. 4743–4751, <http://papers.nips.cc/paper/6581-improved-variational-inference-with-inverse-autoregressive-flow>.
 - [12] X. Chen, D. P. Kingma, T. Salimans, Y. Duan, P. Dhariwal, J. Schulman, I. Sutskever, and P. Abbeel, Variational lossy autoencoder, [arXiv:1611.02731](https://arxiv.org/abs/1611.02731).
 - [13] A. Paszke *et al.*, Pytorch: An imperative style, high-performance deep learning library, in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché Buc, E. Fox, and R. Garnett (Curran Associates, Inc., New York, 2019), pp. 8024–8035, <http://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high->.
 - [14] E. Bingham, J. P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. A. Szerlip, P. Horsfall, and N. D. Goodman, Pyro: Deep universal probabilistic programming, *J. Mach. Learn. Res.* **20**, 1 (2019), <https://www.jmlr.org/papers/v20/18-403.html>.
 - [15] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, emcee: The MCMC Hammer, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
 - [16] S. R. Hinton, Chain Consumer, *J. Open Source Software* **1**, 00045 (2016).
 - [17] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), A guide to LIGO-Virgo detector noise and extraction of transient gravitational-wave signals, *Classical Quantum Gravity* **37**, 055002 (2020).

- [18] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016).
- [19] M. Germain, K. Gregor, I. Murray, and H. Larochelle, MADE: Masked autoencoder for distribution estimation, in *International Conference on Machine Learning* (Proceedings of Machine Learning Research, Lille, 2015), pp. 881–889, <http://proceedings.mlr.press/v37/germain15.html>.
- [20] S. Ioffe and C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shift, in *International Conference on Machine Learning* (Proceedings of Machine Learning Research, Lille, 2015), pp. 448–456, <http://proceedings.mlr.press/v37/ioffe15.html>.
- [21] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, WaveNet: A generative model for raw audio, [arXiv:1609.03499](https://arxiv.org/abs/1609.03499).
- [22] A. Van den Oord, N. Kalchbrenner, L. Espeholt, O. Vinyals, A. Graves *et al.*, Conditional image generation with pixelcnn decoders, in *Advances in Neural Information Processing Systems 29* (Curran Associates, Inc., 2016), pp. 4790–4798, <http://papers.nips.cc/paper/6527-conditional-image-generation-with-pixelcnn-decoders>.
- [23] S.R. Bowman, L. Vilnis, O. Vinyals, A.M. Dai, R. Jozefowicz, and S. Bengio, Generating sentences from a continuous space, [arXiv:1511.06349](https://arxiv.org/abs/1511.06349).
- [24] J. Aasi *et al.* (LIGO Scientific Collaboration), Advanced LIGO, *Classical Quantum Gravity* **32**, 074001 (2015).
- [25] Advanced LIGO anticipated sensitivity curves, LIGO Technical Document No. LIGO-T0900288-v2, 2009.
- [26] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. J. Forteza, and A. Bohé, Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era, *Phys. Rev. D* **93**, 044007 (2016).
- [27] Y.A. LeCun, L. Bottou, G.B. Orr, and K.-R. Müller, Efficient backprop, in *Neural Networks: Tricks of the Trade*, edited by G. Montavon, G.B. Orr, and K.-R. Müller, Lecture Notes in Computer Science (Springer, Berlin, Heidelberg, 2012), Vol. 7700, https://doi.org/10.1007/978-3-642-35289-8_3.
- [28] D.P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [29] H. Fu, C. Li, X. Liu, J. Gao, A. Celikyilmaz, and L. Carin, Cyclical annealing schedule: A simple approach to mitigating KL vanishing, in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)* (Association for Computational Linguistics, 2019), pp. 240–250, <https://dx.doi.org/10.18653/v1/N19-1021>.
- [30] J. Veitch *et al.*, Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library, *Phys. Rev. D* **91**, 042003 (2015).
- [31] Q. Wang, S.R. Kulkarni, and S. Verdú, Divergence estimation for multidimensional densities via k -nearest-neighbor distances, *IEEE Trans. Inf. Theory* **55**, 2392 (2009).
- [32] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, Weight uncertainty in neural networks, *Proc. Mach. Learn. Res.* **37**, 1613 (2015), <http://proceedings.mlr.press/v37/blundell15.html>.