Inferring parameters for binary black hole mergers using normalizing flows

Stephen R. Green
Max Planck Institute for Gravitational Physics
Am Mühlenberg 1
14476 Potsdam, Germany
stephen.green@aei.mpg.de

Jonathan Gair
Max Planck Institute for Gravitational Physics
Am Mühlenberg 1
14476 Potsdam, Germany
jonathan.gair@aei.mpg.de

Abstract

Since the first detection of gravitational waves five years ago, the LIGO and Virgo observatories have detected dozens of signals from merging compact binaries. Each of these is analyzed using Bayesian inference and a standard sampling algorithm to extract the parameters of the system, assuming a signal model and stationary Gaussian noise. In this work, we describe the application of likelihood-free inference with normalizing flows to perform fast-and-accurate inference on gravitational-wave events. We train a neural-spline flow to model the gravitational-wave posterior over the full 15-dimensional space of system parameters for quasicircular binary black holes, given detector strain data. Training data for this consists of 10^6 simulated waveforms, with noise realizations and extrinsic parameters chosen at train time, effectively providing an infinitely large training dataset. Once the model is trained, we perform inference on data from the first detection, GW150914, and show that we recover results consistent with standard sampling algorithms. We thereby establish deep learning as a tool to confront the growing challenges of gravitational-wave inference.

1 Introduction

The first direct detection of gravitational waves was made in September 2015, when LIGO observed a signal from the merger of two black holes 1.4 billion light-years away [1]. Determination of the parameters of the binary that gave rise to the observed strain data was made possible by having accurate models for the signal and using Bayesian inference techniques. Denoting the parameters of the binary (component masses, spins, sky position, etc.) by θ and the measured strain data by s, the aim of this inference is to sample from the posterior distribution, $\theta \sim p(\theta|s)$. The posterior is calculated using Bayes' theorem,

$$p(\theta|s) = \frac{p(s|\theta)p(\theta)}{p(s)}, \tag{1}$$

where $p(s|\theta)$ is the likelihood, $p(\theta)$ is the prior over parameters, and $p(s)$ is a normalizing factor called the evidence. The likelihood is given by the probability distribution of the residual $n \equiv s - h(\theta)$, once the predicted signal waveform $h(\theta)$ associated to θ is subtracted from s. This residual is assumed to be stationary Gaussian noise with a power spectral density (PSD) estimated for the detectors

Third Workshop on Machine Learning and the Physical Sciences (NeurIPS 2020), Vancouver, Canada.

at the time of the event. With a prior specified, the right hand side of (1) can be evaluated up to normalization, and an algorithm such as Markov Chain Monte Carlo or nested sampling [2, 3] is used to obtain samples from the posterior.

Predictions $h(\theta)$ for gravitational-wave signals come from solving the Einstein field equation of general relativity (plus any equations for matter fields in the case of merging neutron stars). The Einstein equation is a nonlinear partial differential equation for the spacetime metric, and solving it numerically for a binary black hole merger can take weeks. Faster models using perturbation theory and fits to numerical relativity have therefore been developed for use in data analysis (see, e.g., [4, 5]). A standard sampling algorithm, however, requires many millions of likelihood (and waveform) evaluations, so even with fast models, inference for a binary black hole system can take days, and this extends to weeks for binary neutron stars. Especially for binary neutron stars, more rapid inference would be desirable for providing rapid alerts to electromagnetic observers for followup observations.

Since the first detection, the LIGO-Virgo Collaboration has published analyses of 50 binary mergers [6], and they have announced over 20 additional triggers in the second half of the most recent observing run [7], the details of which remain to be published. As the rate of detections grows with increasing detector sensitivity, performing parameter inference on each of these is becoming an increasingly costly task. Costs will only increase further with the the use of newer more accurate waveform models that incorporate more physics [8], which will be needed to reduce systematic errors with sensitive detectors.

There is therefore an urgent need for faster inference methods for gravitational waves. Our aim in this work is to demonstrate that a neural-network conditional density estimator $q(\theta|s)$ can be used to perform fast-and-accurate inference for binary black holes in a realistic setting. We use likelihood-free training to learn an approximation to the true posterior, $q(\theta|s) \approx p(\theta|s)$, for any strain data s consistent with the chosen prior, waveform model, and detector noise PSD. Critically, this requires samples from the prior $\theta \sim p(\theta)$ and simulated strain data $s \sim p(s|\theta)$, but no samples from the posterior. Once trained $q(\theta|s)$ can be rapidly sampled, so the ultimate goal is to amortize training time over many events.

There have been several past efforts along similar lines [9–11], but these have all either used a restricted parameter space, or used restricted approximations for the posterior distribution, or only analyzed simulated data. In this work, which is a workshop version of [12], we use a powerful normalizing flow [13] to produce a sufficiently flexible $q(\theta|s)$ to overcome all of these restrictions: we infer all 15 parameters for binary black holes, with data from multiple detectors, and we demonstrate the approach on real interferometer data from the first detection, GW150914. We show that our results are comparable to those produced using standard samplers.

2 Method

Our aim is to train a neural-network conditional density estimator $q(\theta|s)$ to approximate the gravitational-wave posterior $p(\theta|s)$. To achieve this, $q(\theta|s)$ must have sufficient flexibility to capture the nontrivial distribution over θ, as well as the complicated dependence on s. We use the method of normalizing flows [14, 15].

2.1 Normalizing flows

A normalizing flow $f_s : u \mapsto \theta$ is an invertible mapping with simple Jacobian determinant. Given a "base" distribution $\pi(u)$, the normalizing flow defines a conditional probability distribution via the change of variables rule,

$$q(\theta|s) \equiv \pi(f_s^{-1}(\theta)) \left| \det J_{f_s}^{-1} \right| . \tag{2}$$

The idea is to choose $\pi(u)$ such that it can be easily sampled and its density evaluated; then $q(\theta|s)$ inherits these nice properties. We will adopt the standard choice of taking $\pi(u)$ to be multivariate standard normal (of the same dimension D as of the parameter space θ).

To achieve sufficient flexibility, we use a *neural-spline* normalizing flow [13]. We used the implementation in the `nflows` package [16], and took the specific flow to be a composition of rational-quadratic coupling transforms c_s, following the examples in section 5.1 of [13]. Each c_s holds fixed half of the parameters ($u_{1:d}$), and transforms the remaining ones ($u_{d+1:D}$) elementwise according to a coupling

Table 1: Parameters characterizing quasicircular black hole binaries

Parameter	Description	Prior	Extrinsic
(m_1, m_2)	component masses	$[10\ \mathrm{M}_\odot, 80\ \mathrm{M}_\odot]$, $m_1 \geq m_2$	No
ϕ_c	reference phase	$[0, 2\pi]$	No
$t_{c,\mathrm{geocent}}$	time of coalescence	$[-0.1\ \mathrm{s}, 0.1\ \mathrm{s}]$	Yes
d_L	luminosity distance	$[100\ \mathrm{Mpc}, 1000\ \mathrm{Mpc}]$	Yes
(a_1, a_2)	dimensionless spin magnitudes	$[0, 0.88]$	No
$(\theta_1, \theta_2, \phi_{12}, \phi_{JL})$	spin angles	standard [18]	No
θ_{JN}	inclination relative to line-of-sight	$[0, \pi]$, uniform in sine	No
ψ	polarization angle	$[0, \pi]$	Yes
(α, δ)	sky position	uniform over sky	Yes

law,

$$c_{s,i}(u) = \begin{cases} u_i & \text{if } i \leq d, \\ c_i(u_i; u_{1:d}, s) & \text{if } i > d. \end{cases} \tag{3}$$

The $c_i(u_i; u_{1:d}, s)$ are rational-quadratic splines, which are analytically differentiable and invertible with respect to u_i; hence they give rise to a normalizing flow. The composition of coupling transforms, with permutations of the components of u in between, gives rise to a powerful flow. Parameters defining the splines are output from residual neural networks [17], which take as input $u_{1:d}$ and s.

2.2 Training

We train the network to minimize the expected value (over s) of the cross-entropy between the true and model distributions,

$$L = -\int \mathrm{d}s\mathrm{d}\theta\, p(\theta, s) \log q(\theta|s). \tag{4}$$

We evaluate this integral using a Monte Carlo approximation, drawing samples $(\theta^{(i)}, s^{(i)}) \sim p(\theta, s)$ ancestrally, i.e., $\theta^{(i)} \sim p(\theta)$, then $s^{(i)} \sim p(s|\theta^{(i)})$, then

$$L \approx -\frac{1}{N} \sum_{i=1}^{N} \log q(\theta^{(i)}|s^{(i)}), \tag{5}$$

where N is the number of samples in the minibatch. This training approach, called likelihood-free or simulation-based inference, requires no samples from the posterior $p(\theta|s)$, which would be computationally prohibitive. Instead, it merely requires the ability to generate artificial strain data.

2.2.1 Prior

We perform inference over the full 15-dimensional space of binary black hole parameters. These are listed in table 1, along with their prior ranges. (Priors are uniform unless otherwise indicated.) Priors have been chosen to be uninformative and to cover a large region of parameter space that includes GW150914.

2.2.2 Strain data

Given parameters $\theta^{(i)}$, we use the IMRPhenomPv2 frequency-domain waveform model for precessing quasicircular binaries [19, 4, 20] to generate an associated signal $h^{(i)}$. We then add a noise realization $n^{(i)} \sim p_{S_n}(n)$, drawn from a stationary Gaussian distribution with a PSD estimated in the two LIGO detectors in the 1024 s before the event. This is a reasonable approximation to the detector noise, although in the future it would be interesting to study deviations from the stationary Gaussian idealization. Strain data sets s are initially 8 s long and cover a frequency range between 20 Hz and 1024 Hz, but we use a singular value decomposition to compress these to a 100-component reduced-order representation in each detector. These are then whitened and standardized before being fed into the network.

We use a training set of 10^6 elements, which we split into 90% training and 10% test data. However, since it is very fast to generate noise realizations, we do this at train time, effectively enlarging the

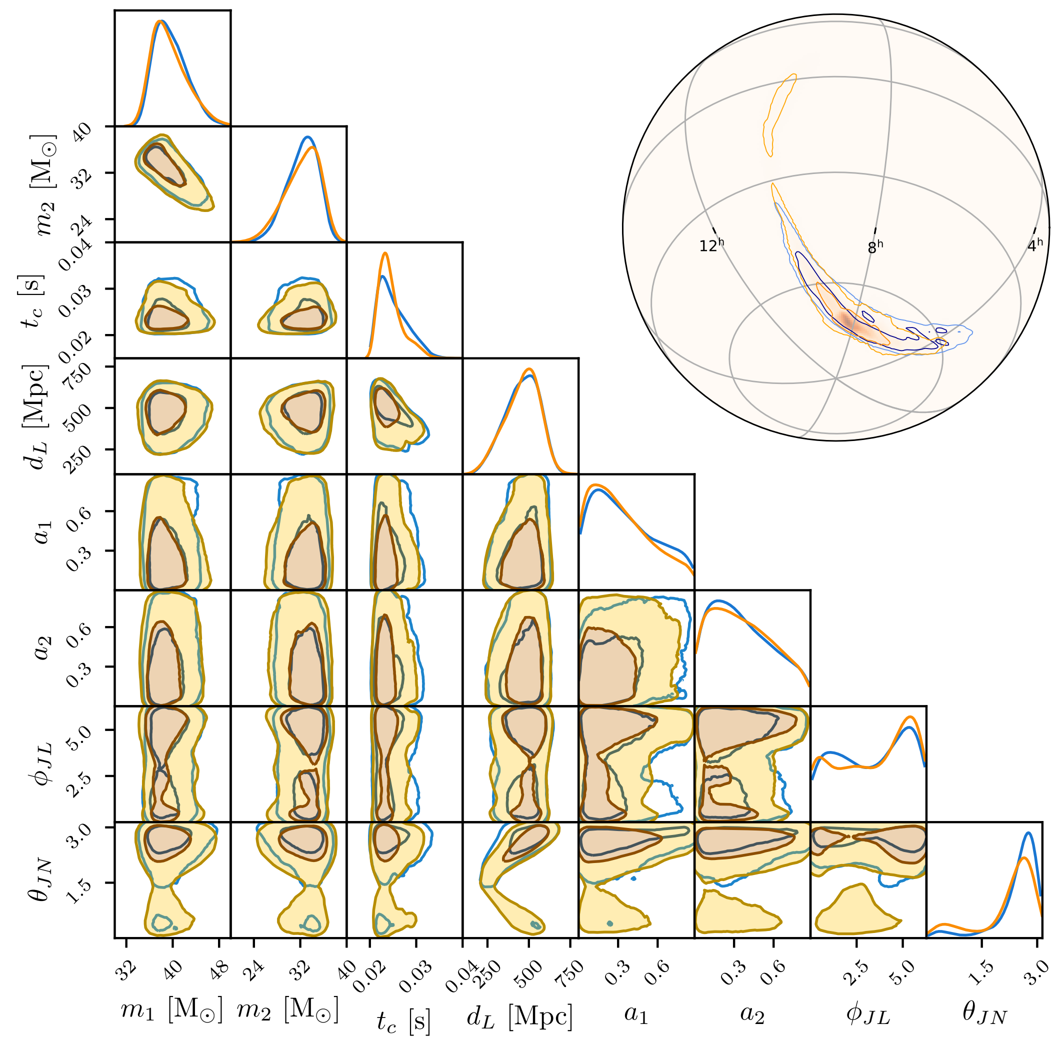

Figure 1: Marginalized one- and two- dimensional posterior distributions over a subset of parameters, comparing the normalizing flow (orange) and `bilby dynesty` (blue). Contours represent 50% and 90% credible regions. The inset shows the sky position. See [12] for a posterior that includes several additional parameters.

size of the training set. In addition, some of the parameters (those labeled "extrinsic" in table 1) are chosen at train time, since they induce simple transformations of the waveforms.

3 Results

We trained for 500 epochs with a batch size of 512. We found that performance was best with 15 coupling transforms, with spline functions given by 9 knots (see [13]). Each neural network was made up of 10 fully-connected residual blocks of 2 layers each, with 512 hidden units. These large residual networks were needed to extract sufficient information from the strain data. With this configuration, training on an NVIDIA Quadro P4000 GPU took approximately 6 days.

After training the network, we used it to perform inference on the actual strain data for GW150914. To do this, we took an 8 s stretch of detector data that included the signal, expressed it in the reduced-order representation, and passed this as context to the flow. The model then produced posterior samples at a rate of 5,000 per second. As a benchmark, we analyzed the same data using the `dynesty` sampler [21] in the `bilby` Bayesian inference library for gravitational-wave data analysis [22, 3].

In figure 1 we present our main result, comparing the two samplers. For the most part, the two posteriors are in very close agreement, with the neural network having larger secondary modes in the sky position and the inclination.

In addition to analyzing GW150914, we performed a series of injection tests with artificial data to ensure good performance throughout parameter space. For each one-dimensional marginalized posterior, we compared the percentile values of the true parameters against a uniform distribution using a Kolmogorov-Smirnov test, and showed they were appropriately distributed (see [12]).

4 Conclusions

In this work, we demonstrated for the first time that deep learning can be used to infer all 15 binary black hole parameters from real gravitational-wave strain data, including accurate estimates of uncertainties. Rapid parameter estimation is critical for multimessenger followup and for confronting the growing number of gravitational-wave detections.

Once trained, our neural conditional density estimator enables inference in seconds. There have been several other applications of deep-learning likelihood-free inference methods to gravitational waves, including the use of gaussian-mixture and histogram models [10], and a conditional variational autoencoder [9]. Although these methods enable similarly fast inference, they either have limited ability to represent complicated posterior distributions, or they are limited to low numbers of dimensions, or they minimize an upper bound on the cross-entropy loss (in the case of the variational autoencoder). The work presented here using a powerful normalizing flow is so-far the only demonstration of accurate inference of all 15 parameters from real data.

The neural density estimators we described are at present tuned to the detector noise characteristics at the time of the event. Although the detectors are mostly stable during an observing run, the noise PSDs do vary slightly from event to event. To fully amortize training costs over every event it will therefore be necessary to also condition the model on this information; alternatively, the network could learn the noise properties directly from the data. The former approach is closest to current analysis methods, and work is underway to incorporate this.

Broader Impact

Since we are analyzing data from gravitational-wave observatories, there are no ethical implications of our work. One of the main long-term aims of our work is to develop deep-learning tools as viable alternatives to standard data analysis methods for gravitational-wave data analysis. These standard approaches are very computationally costly and use a significant amount of electricity. If the approaches we are developing become widespread and reduce electricity usage, then this would have a positive impact on the Earth's climate.

References

[1] B.P. Abbott et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.*, 116(6):061102, 2016, 1602.03837.

[2] J. Veitch et al. Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library. *Phys. Rev.*, D91(4):042003, 2015, 1409.7215.

[3] I.M. Romero-Shaw et al. Bayesian inference for compact binary coalescences with BILBY: Validation and application to the first LIGO–Virgo gravitational-wave transient catalogue. 6 2020, 2006.00714.

[4] Sebastian Khan, Sascha Husa, Mark Hannam, Frank Ohme, Michael Pürrer, Xisco Jimńez Forteza, and Alejandro Bohé. Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era. *Phys. Rev.*, D93(4):044007, 2016, 1508.07253.

[5] Alejandro Bohé et al. Improved effective-one-body model of spinning, nonprecessing binary black holes for the era of gravitational-wave astrophysics with advanced detectors. *Phys. Rev. D*, 95(4):044028, 2017, 1611.03703.

[6] B. P. Abbott et al. GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run. 2020, 2010.14527.

[7] Gravitational-wave candidate event database. `https://gracedb.ligo.org/superevents/public/O3/`. Accessed: 2020-07-22.

[8] Serguei Ossokine et al. Multipolar Effective-One-Body Waveforms for Precessing Binary Black Holes: Construction and Validation. 4 2020, 2004.09442.

[9] Hunter Gabbard, Chris Messenger, Ik Siong Heng, Francesco Tonolini, and Roderick Murray-Smith. Bayesian parameter estimation using conditional variational autoencoders for gravitational-wave astronomy, 2019, 1909.06296.

[10] Alvin J. K. Chua and Michele Vallisneri. Learning Bayesian posteriors with neural networks for gravitational-wave inference. *Phys. Rev. Lett.*, 124(4):041102, 2020, 1909.05966.

[11] Stephen R. Green, Christine Simpson, and Jonathan Gair. Gravitational-wave parameter estimation with autoregressive neural network flows. *Phys. Rev. D*, 102(10):104057, 2020, 2002.07656.

[12] Stephen R. Green and Jonathan Gair. Complete parameter inference for GW150914 using deep learning. 8 2020, 2008.03312.

[13] Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. Neural spline flows, 2019, 1906.04032.

[14] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pages 1530–1538, 2015, 1505.05770.

[15] Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In *Advances in neural information processing systems*, pages 4743–4751, 2016, 1606.04934.

[16] Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. nflows. `https://github.com/bayesiains/nflows`, 2020.

[17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015, 1512.03385.

[18] Benjamin Farr, Evan Ochsner, Will M. Farr, and Richard O'Shaughnessy. A more effective coordinate system for parameter estimation of precessing compact binaries from gravitational waves. *Phys. Rev. D*, 90(2):024018, 2014, 1404.7070.

[19] Mark Hannam, Patricia Schmidt, Alejandro Bohé, Leïla Haegel, Sascha Husa, Frank Ohme, Geraint Pratten, and Michael Pürrer. Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms. *Phys. Rev. Lett.*, 113(15):151101, 2014, 1308.3271.

[20] Alejandro Bohé, Mark Hannam, Sascha Husa, Frank Ohme, Michael Pürrer, and Patricia Schmidt. PhenomPv2 – technical notes for the LAL implementation. *LIGO Technical Document, LIGO-T1500602-v4*, 2016. URL `https://dcc.ligo.org/LIGO-T1500602/public`.

[21] Joshua S Speagle. dynesty: a dynamic nested sampling package for estimating bayesian posteriors and evidences. *Monthly Notices of the Royal Astronomical Society*, 493(3):3132–3158, Feb 2020, 1904.02180. ISSN 1365-2966. URL `http://dx.doi.org/10.1093/mnras/staa278`.

[22] Gregory Ashton et al. BILBY: A user-friendly Bayesian inference library for gravitational-wave astronomy. *Astrophys. J. Suppl.*, 241(2):27, 2019, 1811.02042.