

Complete parameter inference for GW150914 using deep learning

Stephen R. Green* and Jonathan Gair†

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

The LIGO and Virgo gravitational-wave observatories have detected many exciting events over the past five years. As the rate of detections grows with detector sensitivity, this poses a growing computational challenge for data analysis. With this in mind, in this work we apply deep learning techniques to perform fast likelihood-free Bayesian inference for gravitational waves. We train a neural-network conditional density estimator to model posterior probability distributions over the full 15-dimensional space of binary black hole system parameters, given detector strain data from multiple detectors. We use the method of normalizing flows—specifically, a *neural spline* normalizing flow—which allows for rapid sampling and density estimation. Training the network is likelihood-free, requiring samples from the data generative process, but no likelihood evaluations. Through training, the network learns a *global* set of posteriors: it can generate thousands of independent posterior samples per second for any strain data consistent with the prior and detector noise characteristics used for training. By training with the detector noise power spectral density estimated at the time of GW150914, and conditioning on the event strain data, we use the neural network to generate accurate posterior samples consistent with analyses using conventional sampling techniques.

Introduction.—Since the first detection in September 2015 [1], the LIGO/Virgo Collaboration has published observations of gravitational waves from more than a dozen compact binary coalescences [2–5], primarily binary black hole mergers, but also two binary neutron star mergers. In addition, the LIGO/Virgo Collaboration has publicly released around fifty additional triggers [6] of events of interest, the details of which have so far not been published. These observations have had a transformative impact on our understanding of compact objects in the Universe, facilitated by inferring the parameters of the system using accurate physical models of the emitted gravitational waves.

This inference is extremely computationally expensive, as posterior distributions of the parameters are usually obtained using more than one waveform model to probe any potential systematic effects, and using the physically most complete (and thus computationally most costly) waveforms available. LIGO/Virgo currently employ Markov Chain Monte Carlo and nested-sampling algorithms to obtain posterior samples [7, 8]. Run times for single posterior calculations typically take days for binary black hole systems and weeks for binary neutron stars [2, 9]. These long run times will become increasingly problematic as the sensitivity of the instruments improves and event rates reach one per day or higher [10].

There is an urgent need for new approaches that can generate scientific inferences much more rapidly than the existing pipelines [7, 8]. Deep-learning methods are one promising approach to increase the speed of gravitational wave inference by several orders of magnitude, that has been receiving increasing focus in recent years [11]. These techniques attempt to train a neural-network conditional density estimator $q(\theta|s)$ to approximate the Bayesian posterior distribution $p(\theta|s)$ of parameter values θ given detector strain data s . Neural networks typically have

millions of parameters, which are optimized stochastically during training to minimize an appropriate loss function. With a “likelihood-free” training algorithm, it is never necessary to draw samples from the posterior or evaluate a likelihood, rather the procedure is generative and just requires an ability to simulate data sets. Consequently, training can be done in a time comparable to that taken to obtain posterior samples using a standard method. There have been several previous works on this topic, but these either simplified the description of the posterior, e.g., by using a Gaussian approximation [12], or simplified the input, e.g., using a reduced space of parameters and a single detector [13].

In a previous paper [14], we used a type of neural network known as a conditional variational autoencoder (CVAE) [15, 16] combined with normalizing flows [17–20] to learn the posterior distribution for inference of all parameters of an aligned-spin quasi-circular merger observed with a single gravitational-wave detector. With a single detector we could not recover the full set of waveform parameters, and all data sets analyzed were artificially generated with Advanced LIGO design-sensitivity noise [21]. In this Letter, we extend that work into a tool that can, for the first time, be used to analyze real data from the LIGO/Virgo interferometers. We describe a neural network architecture, based on normalizing flows alone, that is able to generate posteriors on the full $D = 15$ dimensional parameter space of quasi-circular binary inspirals, using input data from multiple gravitational-wave detectors. We apply this network to analyze observed interferometer data surrounding the first observed gravitational-wave event, GW150914, and show that we can successfully recover posterior distributions consistent with conventional methods. This is the first demonstration that these methods can be used in a realistic setting to produce fast and accurate scientific inference on real data. This Letter

thus establishes a new benchmark in fast-and-accurate gravitational wave inference, as well as describing methods that could also be applied to other inference problems in experimental physics.

Neural network model.—Our aim is to train a neural conditional density estimator $q(\theta|s)$ to approximate the gravitational-wave posterior $p(\theta|s)$. To this end, $q(\theta|s)$ must have sufficient flexibility to capture the detailed shape of the true posterior over parameters θ , as well as the dependence on the complicated strain data s . We use the method of *normalizing flows*.

A normalizing flow f is an invertible mapping on a sample space with simple Jacobian determinant [17]. For a conditional distribution, the flow must depend on s , so we denote it f_s . The idea is to train the flow so that it maps a simple “base” distribution $\pi(u)$ into the far more complex $q(\theta|s)$. We define the conditional distribution in terms of the flow by

$$q(\theta|s) = \pi(f_s^{-1}(\theta)) \left| \det J_{f_s}^{-1} \right|, \quad (1)$$

which is based on the change of variables rule for probability distributions. The base distribution $\pi(u)$ should be chosen such that it can be easily sampled and its density evaluated; we will always take it to be standard multivariate normal of the same dimension D as the sample space.

By the properties of a normalizing flow, $q(\theta|s)$ inherits the nice properties of the base distribution. Indeed, to draw a sample, one first samples $u \sim \pi(u)$, and then sets $\theta = f_s(u)$; it follows that $\theta \sim q(\theta|s)$. To evaluate the conditional density, one uses (1); the right hand side may be evaluated by the defining properties that f_s is invertible and has simple Jacobian determinant.

Normalizing flows are under active development in computer science, and are usually represented by neural networks. Neural networks are very flexible trainable function approximators, so they can give rise to complex conditional densities. Our previous work [14] used a Masked Autoregressive Flow [19] with affine transformations. In the present work, we use a much more powerful flow called a *neural spline* flow [22]. We use directly the original neural spline flow implementation [26], which we illustrate in figure 1. We now give a brief summary.

The flow is a composition of “coupling transforms” $c_s(u)$, each of which transform elementwise half of the parameters (say, $u_{d+1:D}$) conditional upon the other half ($u_{1:d}$) as well as the strain data s [27], i.e.,

$$c_{s,i}(u) = \begin{cases} u_i & \text{if } i \leq d, \\ c_i(u_i; u_{1:d}, s) & \text{if } i > d. \end{cases} \quad (2)$$

If c_i is invertible and differentiable with respect to u_i , then it follows immediately that the coupling transform is a normalizing flow. By composing n_{flows} of these transforms, and permuting the indices of u in between, a much more flexible flow can be obtained.

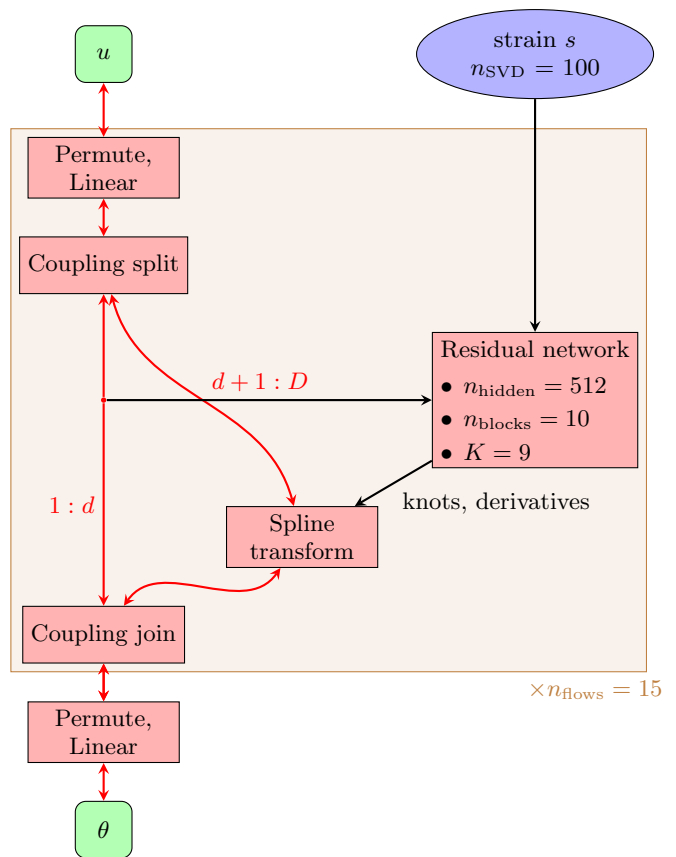


FIG. 1. Overall structure of the normalizing flow [22] from the base space (u) to the parameter space (θ), with optimal hyperparameter choices indicated. Red connections are invertible. The residual network is made up of n_{blocks} residual blocks, each with two fully-connected hidden layers of n_{hidden} units. Prior to each linear transformation [23], we inserted batch normalization layers to speed training [24] and Exponential Linear Units for nonlinearity [25]. Each block is also conditioned on the strain data s .

The neural spline coupling transform [22] takes each c_i to be a monotonically-increasing piecewise function, defined by a set of knots $\{(u_i^{(k)}, c_i^{(k)})\}_{k=0}^K$ and positive-valued derivatives $\{\delta_i^{(k)}\}_{k=0}^K$, between which are interpolated rational-quadratic (RQ) functions. The knots and derivatives are output from a residual neural network [28], which takes as input $u_{1:d}$ and s ; details of the network are given in figure 1. The RQ spline is differentiable and has analytic inverse, so it satisfies the properties of a coupling transform.

Training.—The conditional density estimator $q(\theta|s)$ must be trained to approximate as closely as possible the gravitational-wave posterior $p(\theta|s)$. We do this by tuning the neural network parameters to minimize a loss function, the expected value (over s) of the cross-entropy

between the true and model distributions,

$$\begin{aligned} L &= - \int ds p(s) \int d\theta p(\theta|s) \log q(\theta|s) \\ &= - \int d\theta p(\theta) \int ds p(s|\theta) \log q(\theta|s). \end{aligned} \quad (3)$$

On the second line we used Bayes’ theorem to express L in a form that involves an integral over the likelihood rather than the posterior; this is a key simplification which means posterior samples are not needed for training. We evaluate the integral (3) on a minibatch of training data with a Monte Carlo approximation,

$$L \approx -\frac{1}{N} \sum_{i=1}^N \log q(\theta^{(i)}|s^{(i)}), \quad (4)$$

where $\theta^{(i)} \sim p(\theta)$, $s^{(i)} \sim p(s|\theta^{(i)})$, and N is the number of samples in the minibatch. We then use backpropagation (the chain rule) to compute the gradient with respect to network parameters, and minimize L stochastically on minibatches using the Adam optimizer [29].

To obtain a training pair $(\theta^{(i)}, s^{(i)})$, we draw $\theta^{(i)}$ from the prior, we generate a waveform $h(\theta^{(i)})$, and then we add a noise realization to obtain $s^{(i)}$. Waveform generation is too costly to perform in real time during training, so we adopt a hybrid approach: we sample “intrinsic” parameters in advance and save associated waveform polarizations $h_{+, \times}^{(i)}$; at train time we sample “extrinsic” parameters, project onto detectors, and add noise. We used 10^6 sets of intrinsic parameters, which was sufficient to avoid overfitting.

Prior.—We perform inference over the full $D = 15$ dimensional set of precessing quasi-circular binary black hole parameters: detector-frame masses m_i ($i = 1, 2$), reference phase ϕ_c , time of coalescence $t_{c, \text{geocent}}$, luminosity distance d_L , spin magnitudes a_i , spin angles $(\theta_i, \phi_{12}, \phi_{JL})$ [30], inclination angle θ_{JN} , polarization angle ψ , and sky position (α, δ) . Of these, we consider $(m_i, \phi_c, a_i, \theta_i, \phi_{12}, \phi_{JL}, \theta_{JN})$ to be intrinsic, so they are sampled in advance of training. To analyze GW150914, we take a uniform prior over

$$10 M_{\odot} \leq m_i \leq 80 M_{\odot}, \quad (5a)$$

$$100 \text{ Mpc} \leq d_L \leq 1000 \text{ Mpc}, \quad (5b)$$

$$0 \leq a_i \leq 0.88, \quad (5c)$$

$$-0.1 \text{ s} \leq t_{c, \text{geocent}} \leq 0.1 \text{ s}. \quad (5d)$$

The prior is standard over the remaining quantities. We take $t_{c, \text{geocent}} = 0$ to be the trigger time, and we require $m_1 \geq m_2$.

Although a prior uniform in the comoving source frame [8] would be most physical, we adopted a uniform prior over d_L and an upper bound of 1000 Mpc to more uniformly cover the parameter space and improve training. We applied the physical prior in postprocessing by reweighting samples. Also to improve training,

we rescaled all parameters to have zero mean and unit variance.

Strain data.—For likelihood-free training, we require simulated strain data sets $s^{(i)}$ that arise from the data generative process, $s^{(i)} \sim p(s|\theta^{(i)})$. We assume stationary Gaussian noise, so the gravitational-wave likelihood function is known explicitly, but the likelihood-free approach applies even when this is not the case—e.g., in the presence of non-Gaussian noise—as long as it is possible to simulate data.

We generate training waveforms using the `IMRPhenomPv2` frequency-domain precessing model [31–33]. We take a frequency range of [20, 1024] Hz, and a waveform duration of 8 s. We then whiten $h_{+, \times}^{(i)}$ using the noise PSD estimated from 1024 s of detector data prior to the event. Following [12], we compress the whitened waveforms to a reduced-order representation; we use a singular value decomposition (SVD), and keep the first $n_{\text{SVD}} = 100$ components.

At train time, we sample extrinsic parameters and generate detector signals. This requires a trivial rescaling to apply d_L , and linear combinations of $h_{+, \times}^{(i)}$ to project onto the antenna patterns for the two LIGO detectors. To apply time delays in the reduced-basis (RB) representation, we follow the approach of [34, 35] of pre-preparing a grid of time-translation matrix operators that act on vectors of RB coefficients, using cubic interpolation for intermediate times. Since the transformation to RB space is a rotation, we add white noise directly to the RB coefficients of the whitened waveforms to obtain $s^{(i)}$. Finally, data is standardized to have zero mean and unit variance in each component.

Results.—We trained for 500 epochs with a batch size of 512. The initial learning rate was 0.0002, and we used cosine annealing [36] to reduce the learning rate to zero over the course of training. We performed a search over network hyperparameters, and have listed those with best performance (as measured by final validation loss) in figure 1. During training, we reserved 10% of our training set for validation, and found no evidence of overfitting. With an NVIDIA Quadro P4000 GPU, training took ≈ 6 days.

To perform inference on GW150914, we took 8 s of detector data containing the signal and expressed it in the RB representation. We then drew samples from the base space, and applied the normalizing flow conditioned on the strain data to obtain samples from $q(\theta|s)$. This produced samples at a rate of 5,000 per second. We benchmarked these against samples produced by `bilby` [8, 37] with the `dynesty` sampler [38].

Our main result is presented in figure 2, which compares the neural network and `bilby` posteriors. Both distributions are clearly in very close agreement. There are minor differences in the inclination angle θ_{JN} , where the neural network gives more support to the secondary




FIG. 2. 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. Neural network posteriors are constructed from 5×10^4 samples. The inset shows the sky position, with rejection sampling used to obtain unweighted neural network samples.

mode, and the sky position, where the neural network also develops a small secondary mode. With more training or a larger network, we expect even better convergence.

Although our demonstration has focused on GW150914, the neural network has been trained to generate any posterior consistent with the prior and the given noise PSD. To demonstrate this, performed inference on 100 injec-

tions of waveforms drawn from the prior with added noise. Performance is summarized in the P–P plot shown in figure 3. This shows the cumulative distribution function (CDF) of the percentile scores of each injection parameter within the one-dimensional marginalized posteriors. The percentiles should be distributed uniformly between 0 and 1; since the CDF lies close to the diagonal, we conclude




FIG. 3. P–P plot for 100 artificial strain data sets analyzed by the neural network. For each injection and one-dimensional posterior distribution, we compute the percentile value of the injected parameter. The figure shows the CDF of the injections for each parameter, which should lie close to the diagonal if the network is performing properly. KS test p -values are given in the legend.

that the network is properly sampling the posteriors. This is confirmed by a Kolmogorov-Smirnov (KS) test.

In our experiments, we also varied n_{SVD} , and we found slightly reduced performance as this was increased. This indicates that, although with less compression it should be possible to produce tighter posteriors, better network optimization is required to take full advantage. Indeed, subleading SVD elements contain mostly noise, which makes training more difficult.

Conclusions.—In this Letter, we have demonstrated for the first time that deep neural networks can accurately infer all 15 binary black hole parameters from real gravitational-wave strain data. Once the network is trained, inference is extremely fast, producing 5,000 independent samples per second.

Rapid parameter estimation is critical for multimessenger followup and for confronting the expected high rate of future detections. An advantage of likelihood-free methods is that waveform generation is done in advance of training and inference, rather than at sampling time as for conventional methods. Thus, waveform models that include more physics but may be slower to evaluate [39] can be used to analyze data in the same time as faster models.

The network we presented is tuned to a particular noise PSD—in this case, estimated just prior to GW150914. However, the noise characteristics of the LIGO and Virgo

detectors vary from event to event, and ultimately we would like amortize training costs by building a conditional density estimator that can do inference on any event without retraining for each PSD. One approach would be to condition the model on PSD information: during training, waveforms would be whitened with respect to a PSD drawn from a distribution representing the variation in detector noise from event to event, and (a summary of) this PSD information would be passed to the network as additional context. (PSD samples can be obtained from detector data at random times.) For inference, PSD information would then be passed along with the whitened strain data. Similar approaches could also be used to treat non-Gaussian noise artifacts.

In contrast to CVAEs used in past work [13, 14], normalizing flows have the advantage of estimating the density directly, without any need to marginalize over latent variables. This means that the loss function can be taken to be the cross-entropy (4) rather than an upper bound [15, 16]. Moreover, since $q(\theta|s)$ is a normalized probability distribution, the Bayesian evidence can be obtained as a byproduct. The performance we achieved without latent variables in this work was made possible by the use of a more powerful normalizing flow [22] compared to [14]. As new and more powerful normalizing flows are developed by computer scientists in the future, they will be straightforward to deploy to further improve the performance and capabilities of deep learning for gravitational-wave parameter estimation.

We would like to thank C. Simpson for helpful discussions and for performing training runs during the early stages of this work. Our code was implemented in PyTorch [40], and with the neural spline flow implementation of [26]. Plots were produced with matplotlib [41], ChainConsumer [42] and ligo.skymap [43].

* stephen.green@aei.mpg.de

† jonathan.gair@aei.mpg.de

- [1] B.P. Abbott *et al.* (LIGO Scientific, Virgo), “Observation of Gravitational Waves from a Binary Black Hole Merger,” *Phys. Rev. Lett.* **116**, 061102 (2016), arXiv:1602.03837 [gr-qc].
- [2] B. P. Abbott *et al.* (LIGO Scientific, Virgo), “GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs,” (2018), arXiv:1811.12907 [astro-ph.HE].
- [3] R. Abbott *et al.* (LIGO Scientific, Virgo), “GW190412: Observation of a Binary-Black-Hole Coalescence with Asymmetric Masses,” (2020), arXiv:2004.08342 [astro-ph.HE].
- [4] B.P. Abbott *et al.* (LIGO Scientific, Virgo), “GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4M_{\odot}$,” *Astrophys. J. Lett.* **892**, L3 (2020), arXiv:2001.01761 [astro-ph.HE].

- [5] R. Abbott *et al.* (LIGO Scientific, Virgo), “GW190814: Gravitational Waves from the Coalescence of a 23 Solar Mass Black Hole with a 2.6 Solar Mass Compact Object,” *Astrophys. J.* **896**, L44 (2020), arXiv:2006.12611 [astro-ph.HE].
- [6] “Gravitational-wave candidate event database,” <https://gracedb.ligo.org/superevents/public/03/>, accessed: 2020-07-22.
- [7] J. Veitch *et al.*, “Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library,” *Phys. Rev.* **D91**, 042003 (2015), arXiv:1409.7215 [gr-qc].
- [8] 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,” (2020), arXiv:2006.00714 [astro-ph.IM].
- [9] B. P. Abbott *et al.* (LIGO Scientific, Virgo), “Properties of the binary neutron star merger GW170817,” *Phys. Rev.* **X9**, 011001 (2019), arXiv:1805.11579 [gr-qc].
- [10] B.P. Abbott *et al.* (KAGRA, LIGO Scientific, VIRGO), “Prospects for Observing and Localizing Gravitational-Wave Transients with Advanced LIGO, Advanced Virgo and KAGRA,” *Living Rev. Rel.* **21**, 3 (2018), arXiv:1304.0670 [gr-qc].
- [11] Elena Cuoco *et al.*, “Enhancing Gravitational-Wave Science with Machine Learning,” (2020), arXiv:2005.03745 [astro-ph.HE].
- [12] Alvin J. K. Chua and Michele Vallisneri, “Learning Bayesian posteriors with neural networks for gravitational-wave inference,” *Phys. Rev. Lett.* **124**, 041102 (2020), arXiv:1909.05966 [gr-qc].
- [13] 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), arXiv:1909.06296 [astro-ph.IM].
- [14] Stephen R. Green, Christine Simpson, and Jonathan Gair, “Gravitational-wave parameter estimation with autoregressive neural network flows,” (2020), arXiv:2002.07656 [astro-ph.IM].
- [15] Diederik P Kingma and Max Welling, “Auto-encoding variational bayes,” (2013), arXiv:1312.6114 [stat.ML].
- [16] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra, “Stochastic backpropagation and approximate inference in deep generative models,” in *International Conference on Machine Learning* (2014) pp. 1278–1286, 1401.4082 [stat.ML].
- [17] Danilo Rezende and Shakir Mohamed, “Variational inference with normalizing flows,” in *International Conference on Machine Learning* (2015) pp. 1530–1538, 1505.05770 [stat.ML].
- [18] 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* (2016) pp. 4743–4751, arXiv:1606.04934 [cs.LG].
- [19] George Papamakarios, Theo Pavlakou, and Iain Murray, “Masked autoregressive flow for density estimation,” in *Advances in Neural Information Processing Systems* (2017) pp. 2338–2347, arXiv:1705.07057 [stat.ML].
- [20] Xi Chen, Diederik P Kingma, Tim Salimans, Yan Duan, Prafulla Dhariwal, John Schulman, Ilya Sutskever, and Pieter Abbeel, “Variational lossy autoencoder,” (2016), arXiv:1611.02731 [cs.LG].
- [21] “Advanced LIGO anticipated sensitivity curves,” LIGO Technical Document, LIGO-T0900288-v2 (2009).
- [22] Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios, “Neural spline flows,” (2019), arXiv:1906.04032 [stat.ML].
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Identity mappings in deep residual networks,” (2016), arXiv:1603.05027 [cs.CV].
- [24] Sergey Ioffe and Christian Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *International Conference on Machine Learning* (2015) pp. 448–456, arXiv:1502.03167 [cs.LG].
- [25] Djork-Arn Clevert, Thomas Unterthiner, and Sepp Hochreiter, “Fast and accurate deep network learning by exponential linear units (elus),” (2015), arXiv:1511.07289 [cs.LG].
- [26] Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios, “Neural spline flows,” <https://github.com/bayesiains/nsf> (2019).
- [27] Laurent Dinh, David Krueger, and Yoshua Bengio, “Nice: Non-linear independent components estimation,” (2014), arXiv:1410.8516 [cs.LG].
- [28] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” (2015), arXiv:1512.03385 [cs.CV].
- [29] Diederik P. Kingma and Jimmy Ba, “Adam: A Method for Stochastic Optimization,” (2014), arXiv:1412.6980 [cs.LG].
- [30] 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**, 024018 (2014), arXiv:1404.7070 [gr-qc].
- [31] Mark Hannam, Patricia Schmidt, Alejandro Bohé, Lela Haegel, Sascha Husa, Frank Ohme, Geraint Pratten, and Michael Prrer, “Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms,” *Phys. Rev. Lett.* **113**, 151101 (2014), arXiv:1308.3271 [gr-qc].
- [32] Sebastian Khan, Sascha Husa, Mark Hannam, Frank Ohme, Michael Pürerer, Xisco Jiménez 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**, 044007 (2016), arXiv:1508.07253 [gr-qc].
- [33] Alejandro Bohé, Mark Hannam, Sascha Husa, Frank Ohme, Michael Pürerer, and Patricia Schmidt, “PhenomPv2 – technical notes for the LAL implementation,” LIGO Technical Document, LIGO-T1500602-v4 (2016).
- [34] Priscilla Canizares, Scott E. Field, Jonathan R. Gair, and Manuel Tiglio, “Gravitational wave parameter estimation with compressed likelihood evaluations,” *Phys. Rev. D* **87**, 124005 (2013), arXiv:1304.0462 [gr-qc].
- [35] Rory Smith, Scott E. Field, Kent Blackburn, Carl-Johan Haster, Michael Prrer, Vivien Raymond, and Patricia Schmidt, “Fast and accurate inference on gravitational waves from precessing compact binaries,” *Phys. Rev. D* **94**, 044031 (2016), arXiv:1604.08253 [gr-qc].
- [36] Ilya Loshchilov and Frank Hutter, “Sgdr: Stochastic gradient descent with warm restarts,” (2016), arXiv:1608.03983 [cs.LG].
- [37] Gregory Ashton *et al.*, “BILBY: A user-friendly Bayesian inference library for gravitational-wave astronomy,” *Astrophys. J. Suppl.* **241**, 27 (2019), arXiv:1811.02042 [astro-ph.IM].

- [38] Joshua S Speagle, “dynesty: a dynamic nested sampling package for estimating bayesian posteriors and evidences,” *Monthly Notices of the Royal Astronomical Society* **493**, 31323158 (2020), arXiv:1904.02180 [astro-ph.IM].
- [39] Serguei Ossokine *et al.*, “Multipolar Effective-One-Body Waveforms for Precessing Binary Black Holes: Construction and Validation,” (2020), arXiv:2004.09442 [gr-qc].
- [40] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala, “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., 2019) pp. 8024–8035.
- [41] J. D. Hunter, “Matplotlib: A 2d graphics environment,” *Computing in Science & Engineering* **9**, 90–95 (2007).
- [42] S. R. Hinton, “ChainConsumer,” *The Journal of Open Source Software* **1**, 00045 (2016).
- [43] Leo Singer, “ligo.skymap,” <https://lscsoft.docs.ligo.org/ligo.skymap/> (2020).