[LIGO-P2300288-v3]

Rapid neutron star equation of state inference with Normalising Flows

Jordan McGinn,[1] Arunava Mukherjee,[2] Christopher Messenger,[1] Michael J. Williams,[1] and Ik Siong Heng[1]
[1]*SUPA, School of Physics and Astronomy*
University of Glasgow
Glasgow G12 8QQ, United Kingdom
[2]*Saha Institute of Nuclear Physics, 1/AF Bidhannagar, Kolkata-700064, India*
(Dated: February 11, 2022, document in preparation.)

The first detection of gravitational waves from binary neutron stars on the 17th of August, 2017, (GW170817) heralded the arrival of a new messenger for probing neutron star astrophysics. One of the many spectacular outcomes of GW170817 was the first constraints on neutron star equation of state from gravitational wave observations. A significant amount of computing time and resources were used to obtain these first constraints. However, as observations of binary neutron star coalescence become more routine in the coming observing runs, there is a need to significantly improve the analysis speed and flexibility. Here, we present a rapid approach for inferring the neutron star equation of state based on normalising flows. We demonstrate that, using the same input data, our approach, `ASTREOS`, produces consistent with the results presented in Abbott et al. but will only require $\mathcal{O}(0.1)s$ to generate neutron star equation of state curves. Furthermore, `ASTREOS` allows for non-parametric equation of state inference. This rapid analysis will not only feed into neutron star equation of state studies but can potentially feed into future alerts for electromagnetic follow-up observations of binary neutron star observations. Additionally, we discuss the implications this rapid analysis, especially in the context of other rapid Bayesian inference frameworks such as Dingo and VItamin.

Introduction.— On 17 August 2017, during the second advanced detector observing runs, the advanced LIGO [1] and Virgo [2] observatories detected the first gravitational wave (GW) signal from the coalescence of two neutron stars (NSs) [3]. During the 3rd observing run the global network made a second GW observation consistent with a signal from a binary neutron star collision [4]. In addition, a further two signals likely coming from the mergers between pairs of neutron stars and black holes have been reported [5]. These detections have provided a new opportunity to probe matter in extreme conditions such as those in the interior of these stars [6]. To date there have been 90 definitive GW event detections with the vast majority being the merger of binary black holes (BBHs) [7]. However, the effect on matter for merging systems containing one or more NSs are significantly different to those of BBHs as tidal effects have to be taken into account when modelling the waveform.

It is the deformation that each star's gravitational field induces on their partner that accelerates the decay of the inspiral and imprints itself on the emitted GW. GWs provide information on the star's tidal interactions via a neutron star's tidal deformability, Λ. This dimensionless quantity measures the change of a neutron star's quadrupole moment in response to an external tidal field which, in turn, allows us to place effective constraints on the star's equation of state (EOS). As GW detectors steadily improve in sensitivity and therefore increase their rate of detection, we expect to detect many more NS interactions [8] undoubtedly leading to more accurate measurements of the neutron star EOS.

The EOS is assumed to be universal in the sense that there is a single relationship between pressure and energy density within a neutron star that all such stars obey. Each star is free to have its own independent gravitational mass and associated tidal deformability governed by the EOS and the corresponding prior range of allowed masses. Therefore accurate measurements of the neutron star EOS requires the combination of data from multiple neutron star interactions. A fully Bayesian approach was conducted and was shown in the literature. The work was expanded when Markov chain Monte Carlo simulations were used to fit a model that follows the piece-wise polytropic parametrisation of the equation of state.

Recent advances in machine learning (ML) have been successfully applied to GW data analysis [9]. These include parameter estimation [10–13], waveform modelling [14, 15] and searches [16–18]. The motivation behind using ML to either supplement or replace existing algorithms is that, once trained the ML algorithm can run at a significant fraction of the computational cost and/or time. It is also naturally more flexible and can account for hard-to-model or un-modelled components of the analysis.

In this letter, we describe an ML approach to provide posterior distributions describing the neutron star EOS using as input, standard parameter estimation data products (posterior samples) on the gravitational masses and the tidal deformabilities of the components within a binary neutron star (BNS) system. Such inputs are generated as standard as part of the data release associated with detected BNS events, e.g., [19]. The output of our Normalising-Flow[20, 21] based analysis is an ensemble of non-parametric neutron star EOSs describing the relationship between pressure and energy density within neutron stars in addition to correlated samples of central densities/pressures and maximum permitted neutron star densities. Hence, once trained, the Normalising-

Flow acts as a rapid functional generator of plausible neutron star EOSs.

Method.— A Normalising-Flow (or Flow) is a generative ML model that learns to transform the probability density of a set of data to a simpler "latent" distribution via a series of invertible mappings. A normalising flow (NF) has the advantage over other generative models (Generative Adversarial Networks [22], Variational Auto-Encoders [23]) as it predicts the probability density function of the training data explicitly. At the training stage, a conditional Flow model learns to map samples x in the data space $\mathcal{X}$ and y in the conditional data space $\mathcal{Y}$ to a point z in latent space $\mathcal{Z}$ such that the following function holds:

$$p_{\mathcal{X}|\mathcal{Y}}(x|y) = p_{\mathcal{Z}|\mathcal{Y}}(f(x|y)|y)\left|\det\left(\frac{\partial f(x|y)}{\partial x^T}\right)\right|, \quad (1)$$

where f is the bijective mapping $f : \mathcal{X} \to \mathcal{Z}$ and $\partial f(x|y)/\partial x^T$ is the Jacobian of f at x. For the best practical use we choose transforms whose Jacobian determinants are easy to compute. In this work we use a real non-volume preserving (Real NVP) Normalising-Flow [24] which is defined by its use of stacked *coupling transforms*. As the learned transforms are invertible, samples can be drawn from $p_{\mathcal{X}|\mathcal{Y}}$ by sampling from $p_{\mathcal{Z}|\mathcal{Y}}$ and applying the inverse transform f^{-1}. Further details on NFs, their different types and their inner works please see [20, 21].

Our aim is to use a conditional Flow to approximate a function that maps source parameters $y = (m_1, m_2, \Lambda_1, \Lambda_2)$ to an associated set of EOS information x parameterised by a principle component analysis (PCA) decomposition, two values of central density and a maximally allowed central density. We refer to these latter 3 quantities as *auxiliary parameters* and we call our analysis package `ASTREOS`.

The training of `ASTREOS` and indeed most Normalising-Flows involves the process of learning a forward mapping from the training data (the EOS and auxiliary parameters) conditional on labels (the NS component masses and tidal deformabilities) to a zero-mean unit-variance uncorrelated multi-dimensional Gaussian with dimensions equalling that of the training data. Once trained, we can use `ASTREOS` to perform the inverse mapping from a single condition label y and a randomly drawn location z from the latent space distribution, to an EOS and corresponding auxiliary values x. As we can continually draw random points from the latent space to produce EOS data using the same conditional labels, there will naturally be variation in the output EOS data. This is encoded in the Flow output distribution $p_{\mathcal{X}|\mathcal{Y}}(x|y)$. The variation within this distribution is representative of the degeneracy inherent within EOS inference based on single estimates of component masses and tidal deformabilities. A single value of y maps to a distribution of plausible EOSs all consistent with the input conditional data and the prior distribution represented by the training data. We repeat this process over a set of mass and tidal deformability samples drawn from the joint posterior $p(y|h)$, where h is the GW strain data for a particular BNS event. By doing this we are able to marginalise over the correlated uncertainties in y due to the GW detector noise and other correlations between these and other GW parameters. Hence our final result is

$$\begin{aligned} p(x|h) &= \int p_{\mathcal{X}|\mathcal{Y}}(x|y)p(y|h)\,dy \\ &\approx \frac{1}{N}\sum_{j=1}^{N} p_{\mathcal{X}|\mathcal{Y}}(x|y_j)\big|_{y_j\sim p(y|h)} \end{aligned} \quad (2)$$

where N is the number of posterior samples used as input to the trained Flow. In practice, rather than evaluating this function directly, samples of x are drawn from $p(x|h)$ by drawing equal numbers (usually only one) of samples of x from $p_{\mathcal{X}|\mathcal{Y}}(x|y_j)$ for each value y_j.

Equation-of-State data.— We simulate 10^5 phenomenological neutron star EOSs to train the Flow model. To accommodate the analysis within limited computational resources, we generated the EOSs from a 3-piece polytropic neutron star EOS-family widely used in the literature [25]. Each EOS contains a low-density crust as described by the SLy4 EOS but at higher densities behaves as a piece-wise polytrope with transition densities at $10^{14.7}$ g cm^{-3} and 10^{15} g cm^{-3} (CHECK the values!). We empirically choose the polytropic indices and their distributions in such a way that the variation in our EOS training set closely follows the prior data-set used in [3].

Within our training set each EOS consists of energy density pre-computed on a fixed grid of 256 pressure values. As a preprocessing step we remove densities lower than 6^{34} Kg/m^3 as it is only necessary to learn the high density and varying regions. For each EOS there exists a distribution of possible neutron star masses. We define the lower bound of this range to be $1\text{M}_\odot$ up to a maximum possible mass allowed by the EOS. We sample m_1 and m_2 on this range over a uniform prior distribution ensuring $m_1 > m_2$ and in addition, retain the maximum possible mass M_{max} allowed by the EOS. The EOS and masses then define the central densities of each star in the binary, the tidal deformabilities, and the maximum energy density of the EOS.

In order to reduce dimensionality of the problem, ensuring stable and fast training, we use a PCA representation [26] of the energy-density to reduce the EOS energy-density as a function of pressure, to 8 principle components. The data space $\mathcal{X}$ is therefore reduced to 11 dimensions (8 PCA components plus the 3 auxiliary parameters). As a final preprocessing step both EOS and auxiliary components are standardised separately by removing the mean and scaling to have unit variance. Further, the tidal deformability values are represented by their natural logarithms before input to training.

We use an implementation of Normalising-Flows called `GLASFLOWS` [27] based on `NFLOWS` [28] which is written for `PYTORCH` [29]. We train the flow for 5000 epochs with a batch size of 1000 and an initial learning rate of 0.0002

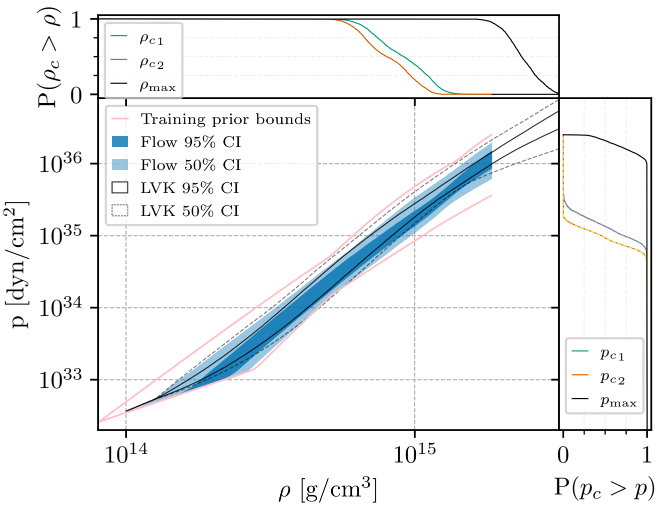

FIG. 1. Generated posterior (blue bands) of pressure p as a function of energy density ρ sampled from `ASTREOS` and posteriors taken from GW170817 analysis (grey lines). Generated central densities and maximum density cumulative distributions from `ASTREOS` and training upper and lower priors (black). Each of the generated data is the result of a learned inverse mapping from a randomly drawn point in a standard Gaussian base distribution and conditional values taken from LIGO analysis. The dark, light and lightest shaded regions correspond to the 95%, 50% confidence intervals respectfully. Top and side panel shows the generated posterior cdfs on $\rho_{c,1}$, $p_{c,1}$ (green), $\rho_{c,2}$, $p_{c,2}$ (orange), ρ_{max}, p_{max} (black). Central pressure cdfs are plotted as dotted lines in the side panel.

which is gradually decayed to zero using cosine annealing together with Adam [30] as the optimiser. There are 5 residual blocks per transform each containing 128 neurons and there are 11 transforms. We use Batch Normalisation between each coupling transform as described in [24] and optimise hyper-parameters using the "weights and bias" package [31]. Data is split between 80% training and 20% validation and on a NVIDIA Tesla V100-PCIE-32GB training requires $\mathcal{O}(12)$ hours.

Results.— We demonstrate the effectiveness of our model by analysing the GW170817 event using posterior samples available from Gravitational Wave Open Science Center [32]. The accompanying paper for this release [6] shows that by explicitly assuming both progenitors were NSs, tighter constraints can be placed on the EOS and mass and tidal deformability parameters. This inherently defines a correlation between tidal deformabilities and assumes that both NSs are in a ground state equilibrium and share the same universal EOS. To perform GW170817 EOS inference using `ASTREOS` we used as input the posterior samples of $m_1, m_2, \Lambda_1, \Lambda_2$ from [19] to represent the conditional y component in the Flow. We draw one latent space sample from $p_{\mathcal{Z}|\mathcal{Y}}$ for each instance of y and apply the Flow to obtain the PCA reduced basis representation of the EOS and each star's associated central densities and EOS maximum permitted density, stored within x. We repeat this process according to Eq. 2 to build up an ensemble of samples of x from the final EOS posterior $p(x|h)$. It takes $\mathcal{O}(0.1)$s to generate and convert 2500 parameter estimation (PE) samples to renormalise EOS curves. Our main GW170817 result is compared to that of [6] in Fig. 1 where we show EOS confidence intervals on pressure as a function of energy-density. We also compare cumulative probability curves for pressure and energy-density for each NS and similar curves for the corresponding maximum allowed pressures and energy-densities of the inferred EOS.

In order to further verify our results we employ a circular argument that tests whether the EOSs and auxiliary parameters generated by the Flow are indeed consistent with the specific $m_1, m_2, \Lambda_1, \Lambda_2$ inputs fed into the Flow. To be clear, each input can map to a family of allowed EOSs and auxiliary parameters, but any such output, when used as input to an EOS solver should return the original specific masses and tidal deformabilities. In performing this test we sample 6 locations in the $m_1, m_2, \Lambda_1, \Lambda_2$ space spanned by the posterior samples from GW170817 [19] specifically ensuring a representative spread in the m_1, m_2 space. For each of these locations we generate 100 instances of EOS*s* and auxiliary parameters using `ASTREOS`. Component masses and tidal deformabilities are then computed for each EOS instance using a Tolman-Oppenheimer-Volkoff (TOV) solver. Figure 2 shows the joint distributions of recovered masses and tidal deformabilities in reference to the original values. Typical variation seen across the GW170817 posterior space is $\sim \pm 0.1\mathrm{M}_\odot$ and $\sim \pm 10$ in component mass and tidal deformability respectively.

Conclusions.— In this Letter we have for the first time demonstrated that neural networks can accurately infer neutron star EOS curves conditioned on prior measurements of component masses and tidal deformabilities. The analysis also provides estimates of the central densities of the stars and maximally allowed energy density of the EOS. Once trained, the network can rapidly generated these curves at 50000 generations per second meaning that EOS posteriors can be built almost immediately after a GW PE run. As the training regime is completely model independent and solely based on expansive training data priors, this model is flexible enough to run on various events without the costly need to retrain. Further to this, the analysis can be easily modified to handle neutron star-black hole (NSBH) systems [5, 7]. Due to the quick and flexible generative properties of this analysis, further constraints on a universal EOS can be found by combining independent BNS and NSBH events hierarchically. For super low latency results, this work can be coupled to existing rapid ML PE codes [10, 12].

This document has LIGO DCC preprint number: `LIGO-P2300288-v3`.

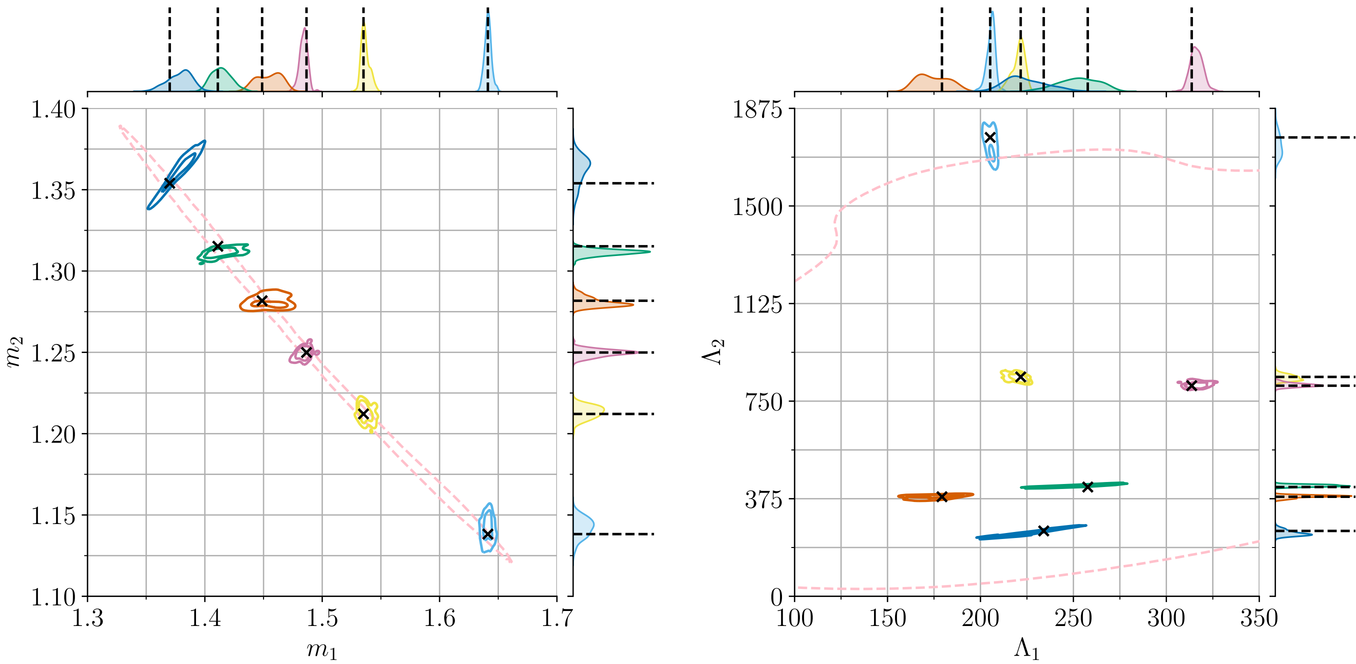

FIG. 2. Recovered component masses and tidal deformaibilites from integrating the TOV equations on flow generated EOS. We take six sets of masses and tidal deformabilites from GW170817 posteriors (grey crosses/dotted lines) and use them to generate 100 EOS's each with `ASTREOS`. These curves are then solved with the TOV equations and the outputs plotted as probability distributions over mass and tidal deformability. Each set is assigned a different colour and we plot the 95% CI over mass and tidal deformability from GW170817 posterior (dotted pink).

[1] J. Aasi *et al.* (LIGO Scientific), Advanced LIGO, Class. Quant. Grav. **32**, 074001 (2015), arXiv:1411.4547 [gr-qc].

[2] F. Acernese *et al.* (VIRGO), Advanced Virgo: a second-generation interferometric gravitational wave detector, Class. Quant. Grav. **32**, 024001 (2015), arXiv:1408.3978 [gr-qc].

[3] B. Abbott and et al., Gw170817: Observation of gravitational waves from a binary neutron star inspiral, Physical Review Letters **119**, 10.1103/physrevlett.119.161101 (2017).

[4] B. P. Abbott and et al., GW190425: Observation of a compact binary coalescence with total mass ~ 3.4 m $\odot$, The Astrophysical Journal **892**, L3 (2020).

[5] R. Abbott and et al., Observation of gravitational waves from two neutron star–black hole coalescences, The Astrophysical Journal Letters **915**, L5 (2021).

[6] B. P. Abbott and et al. (The LIGO Scientific Collaboration and the Virgo Collaboration), Gw170817: Measurements of neutron star radii and equation of state, Phys. Rev. Lett. **121**, 161101 (2018).

[7] The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration, *et al.*, GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run, arXiv e-prints , arXiv:2111.03606 (2021), arXiv:2111.03606 [gr-qc].

[8] The LIGO Scientific Collaboration, the Virgo Collaboration, *et al.*, Prospects for Observing and Localizing Gravitational-Wave Transients with Advanced LIGO and Advanced Virgo, Living Reviews in Relativity **19**, 1 (2016).

[9] E. Cuoco and et al., Enhancing gravitational-wave science with machine learning, Mach. Learn.: Sci. Technol. 2 011002 **2**, 011002 (2020).

[10] H. Gabbard and et al., Bayesian parameter estimation using conditional variational autoencoders for gravitational-wave astronomy (2020), arXiv:1909.06296 [astro-ph.IM].

[11] A. J. K. Chua and M. Vallisneri, Learning bayesian posteriors with neural networks for gravitational-wave inference, Phys. Rev. Lett. **124**, 041102 (2020).

[12] S. R. Green and et al., Gravitational-wave parameter estimation with autoregressive neural network flows, Phys. Rev. D **102**, 104057 (2020).

[13] M. J. Williams and et al., Nested sampling with normalizing flows for gravitational-wave inference, Phys. Rev. D **103**, 103006 (2021).

[14] J. McGinn, C. Messenger, M. J. Williams, and I. S. Heng, Generalised gravitational wave burst generation with generative adversarial networks, Classical and Quantum Gravity **38**, 155005 (2021).

[15] D. Williams, I. S. Heng, J. Gair, J. A. Clark, and B. Khamesra, Precessing numerical relativity waveform surrogate model for binary black holes: A gaussian process regression approach, Phys. Rev. D **101**, 063011 (2020).

[16] D. George and E. A. Huerta, Deep Learning for Real-time Gravitational Wave Detection and Parameter Estimation with LIGO Data, arXiv e-prints , arXiv:1711.07966 (2017), arXiv:1711.07966 [gr-qc].

[17] H. Gabbard, M. Williams, F. Hayes, and C. Messenger, Matching matched filtering with deep networks for gravitational-wave astronomy, Phys. Rev. Lett. **120**, 141103 (2018).

[18] J. Bayley, C. Messenger, and G. Woan, Robust machine learning algorithm to search for continuous gravitational waves, Phys. Rev. D **102**, 083024 (2020).

[19] https://dcc.ligo.org/ligo-p1800115/public.

[20] I. Kobyzev, S. J. Prince, and M. A. Brubaker, Normalizing flows: An introduction and review of current methods, IEEE Transactions on Pattern Analysis and Machine Intelligence **43**, 3964–3979 (2021).

[21] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, Normalizing flows for probabilistic modeling and inference (2021), arXiv:1912.02762 [stat.ML].

[22] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, Generative adversarial nets, in *Advances in Neural Information Processing Systems*, Vol. 27, edited by Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger (Curran Associates, Inc., 2014).

[23] D. P. Kingma and M. Welling, Auto-encoding variational bayes (2014), arXiv:1312.6114 [stat.ML].

[24] L. Dinh, J. Sohl-Dickstein, and S. Bengio, Density estimation using real nvp (2017), arXiv:1605.08803 [cs.LG].

[25] J. S. Read, B. D. Lackey, B. J. Owen, and J. L. Friedman, Constraints on a phenomenologically parametrized neutron-star equation of state, Phys. Rev. D **79**, 124032 (2009), arXiv:0812.2163 [astro-ph].

[26] I. Jolliffe, *Principal component analysis* (Springer Verlag, New York, 2002).

[27] M. J. Williams, J. McGinn, and F. Stachurski, Glasflows, `https://github.com/igr-ml/glasflow` (2021).

[28] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, nflows: normalizing flows in PyTorch (2020).

[29] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, Pytorch: An imperative style, high-performance deep learning library, in *Advances in Neural Information Processing Systems 32* (Curran Associates, Inc., 2019) pp. 8024–8035.

[30] D. P. Kingma and J. Ba, Adam: A Method for Stochastic Optimization, arXiv e-prints , arXiv:1412.6980 (2014), arXiv:1412.6980 [cs.LG].

[31] L. Biewald, Experiment tracking with weights and biases (2020), software available from wandb.com.

[32] Gravitational wave open science center (gwosc).