

SC17: G: Deep Learning for Time-series Signal Processing for Real-time Gravitational Wave Detection and Parameter Estimation: Results with Real LIGO Data

Daniel George^{1,2} and E. A. Huerta²

¹*Department of Astronomy, University of Illinois at Urbana-Champaign, Urbana, Illinois, 61801*

²*NCSA, University of Illinois at Urbana-Champaign, Urbana, Illinois, 61801*

The recent Nobel-prize-winning detections of gravitational waves from merging black holes and the subsequent detection of the collision of two neutron stars in coincidence with electromagnetic observations have inaugurated a new era of multimessenger astrophysics. To enhance this emergent science, we propose the use of deep learning with a system of 1D convolutional neural networks, that take time-series inputs, for classification and regression with a novel curriculum learning scheme and transfer learning technique. We demonstrate how this method can be applied for rapid *detection* and *parameter estimation* of gravitational waves from binary black hole mergers using continuous data streams from multiple LIGO detectors. After training with data derived from high-performance physics simulations, we demonstrate for the first time that machine learning can directly detect and estimate the true parameters of *real* events observed by LIGO. Our results show that deep learning can achieve similar sensitivities and significantly lower errors compared to state-of-the-art matched-filtering searches while being far more computationally efficient and more resilient to non-Gaussian anomalies, allowing real-time processing of weak time-series signals in non-stationary non-Gaussian noise with minimal resources, and also enables the detection of new classes of gravitational wave sources that may go unnoticed with existing detection algorithms. This unified framework for data analysis is ideally suited to enable coincident detection campaigns of gravitational wave sources and their electromagnetic counterparts in real-time.

I. INTRODUCTION

The first detection (GW150914) of gravitational waves (GWs), from the merger of binary black holes (BBHs), with the advanced Laser Interferometer Gravitational-wave Observatory (LIGO) has set in motion a scientific revolution [1] leading to the Nobel prize in Physics in 2017. This was followed by subsequent groundbreaking discoveries [2–4] including the first detection of GWs from the merger of two neutron stars in conjunction with follow-up observations across the electromagnetic (EM) spectrum [5]. This event has finally solved long-standing mysteries proving that neutron star mergers are the creators of heavy metals such as Gold and Platinum and are the central engines of short gamma ray bursts [6, 7].

Matched-filtering, the most sensitive state-of-the-art GW detection algorithm used by LIGO, currently targets only a 4D parameter space (compact binary sources with spin-aligned components on quasi-circular orbits) [8]—a subset of the 9D parameter space available to GW detectors [9, 10]. Extending these template-matching searches to target spin-precessing, quasi-circular or eccentric BBHs is computationally prohibitive [9, 11].

Based on these critical issues, we need a new paradigm to overcome the limitations and computational challenges of existing GW detection algorithms. An ideal candidate would be the rapidly advancing field called Deep Learning, which is a highly scalable machine learning technique that can learn directly from raw data, without any manual feature engineering, by using deep hi-

erarchical layers of “artificial neurons”, called neural networks, in combination with optimization techniques based on back-propagation and gradient descent [12, 13]. Deep learning, especially with the aid of GPU computing, has recently achieved immense success in both commercial applications and artificial intelligence (AI) research [12, 14–16] and has also been applied in astrophysics [17–21]

In this article, we propose for the first time a method based on deep learning for gravitational wave detection. This method, *Deep Filtering*, employs a system of two deep convolution neural networks (CNNs [22]) that directly take time-series inputs for both classification and regression with highly noisy time-series data streams to detect weak GW signals and estimate the parameters of their source in real-time. We demonstrate that deep learning can be used for both signal detection and multiple-parameter estimation directly from extremely weak time-series signals embedded in highly non-Gaussian and non-stationary noise, once trained with some templates of the expected signals. Our results show that deep CNNs achieve performance comparable to matched-filtering methods, while being several orders of magnitude faster and far more resilient to transient noise artifacts such as glitches. We also illustrate how *Deep Filtering* can deal with data streams of arbitrary length from multiple detectors. Most importantly, this article shows for the very *first* time that machine learning can successfully detect and recover true parameters of *real* GW signals observed by LIGO. Furthermore, we show that after a single training process, *Deep*

Filtering can automatically generalize to noise having new Power Spectral Densities (PSDs) from different LIGO events, without re-training.

Our results indicate that `Deep Filtering` can interpolate between elements in the template bank, generalize to new classes of signals beyond the training data, and, surprisingly, detect GW signals and measure their parameters even when they are contaminated by glitches. We present experiments demonstrating the robustness of `Deep Filtering` in the presence of transient non-Gaussian anomalies (glitches). Deep learning, in principle, can learn characteristics of noise in the LIGO detectors and develop better strategies than matched-filtering, which is known to be optimal only for Gaussian noise. Since all the intensive computation is diverted to the one-time training stage of the CNNs, template banks of practically any size may be used for training after which continuous data streams can be analyzed in real-time with a single CPU, while very intensive searches can be rapidly carried out using a single GPU. The initial estimates provided by `Deep Filtering` may be used to instantly narrow down the parameter space of GW detections, which can then be followed up with existing pipelines using a few templates around the predicted parameters to obtain more informative parameter estimates, thus accelerating GW analysis with minimal computational resources across the full range of signals.

II. METHOD

`Deep Filtering` consists of a classifier CNN and a predictor CNN. The classifier has an additional softmax layer which returns probabilities for True or False depending on whether a signal is present. The classifier is first applied to the continuous data stream via a sliding window. If the classifier returns higher probability for True, the predictor is applied to the same input to determine the parameters of the source. In a multi-detector scenario, `Deep Filtering` may be applied separately to each data stream and the coincidence of detections with similar parameters would strengthen the confidence of a true detection, which can be quickly verified by matched-filtering with the predicted templates.

The datasets of waveform templates used to train our CNNs were obtained using the effective-one-body (EOB) code [23] which is calibrated with massively parallel numerical relativity simulations solving Einstein's Field Equations [24]. Our training set contained about 2500 templates, with BBHs component masses sampled in the range $5M_{\odot}$ to $75M_{\odot}$ in steps of $1M_{\odot}$. The testing dataset also contained approximately 2500 templates with intermediate component masses separated from the training set by $0.5M_{\odot}$ each. Subsequently, we produced copies of each signal by shifting the location of their

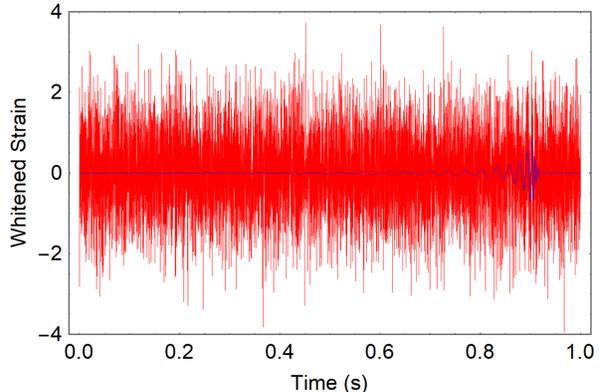


FIG. 1. **Sample signal injected into real LIGO noise.** The red time-series is an example of the input to our algorithm. It contains a hidden BBH GW signal (blue) from our test set which was superimposed in real LIGO noise from the test set and whitened. For this injection, the optimal matched-filter SNR = 7.5 (peak power of this signal is 0.65 times the power of background noise). The component masses of the merging BHs are $57M_{\odot}$ and $33M_{\odot}$. The presence of this signal was detected directly from the (red) time-series input with over 99% sensitivity and the source's parameters were estimated with a mean relative error less than 10%.

peaks randomly within the final 0.2 seconds to make the CNNs more resilient to time translations. This means that in practice, our algorithm will be applied to the continuous data stream using a 1-second sliding window with offsets of 0.2 seconds.

We obtained real LIGO data from the LIGO Open Science Center (LOSC) around the first 3 GW events, namely, GW150914, LVT151012, and GW151226. Each event contained 4096 seconds of real data from each detector. We used noise sampled from GW151226 and LVT151012 for training and validation of our model and noise from GW150914 was used for testing. These tests ensure that our model is able to generalize to different noise distributions, also in the presence of transient glitches, since it is well known that the PSD of LIGO is highly non-stationary, varying widely with time. Therefore, if `Deep Filtering` performs well on these test sets, it would also perform well on data from future time periods, without being re-trained.

Next, we superimposed different realizations of noise randomly sampled from the training set of real LIGO noise from the two different events (GW151226 and LVT151012) and injected signals over multiple iterations, thus amplifying the size of the training datasets. The power of the noise was adjusted according to the desired optimal matched-filter Signal-to-Noise Ratio (SNR) for each training round. The inputs were then whitened with the average PSD of the real noise measured at that time-period. We also scaled and mixed dif-

ferent samples of LIGO noise together to artificially produce more training data and various levels of Gaussian noise was also added to augment the training process. However, the testing results were measured using only pure LIGO noise not used in training with true GW signals or with signals injected from the unaltered test sets.

There were 4 convolution layers with the filter sizes to 64, 128, 256, and 512 respectively and 2 fully connected layers with sizes 128 and 64. The standard ReLU activation function, $\max(0, x)$, was used throughout as the non-linearity between layers. We used kernel sizes of 16, 16, 16, and 32 for the convolutional layers and 4 for all the (max) pooling layers. Stride was chosen to be 1 for all the convolution layers and 4 for all the pooling layers. We observed that using dilations [25] of 1, 2, 2, and 2 in the corresponding convolution layers improved the performance. This CNN architectures was optimized to deal with only Gaussian noise having a flat PSD. However, we later found that this model also obtained the best performance with noise having the colored PSD of LIGO, among all the models we tested. This indicates that our CNNs will be robust to a wide range of noise distributions. Furthermore, pre-training the CNNs on Gaussian noise (transfer learning) before fine-tuning on the limited amount of real noise prevented over-fitting, i.e., memorizing only the training data without generalizing to new inputs. The learning algorithm was ADAM [26].

While training, we developed a curriculum learning strategy to improve the performance and reduce training times of the CNNs while retaining performance at very high SNR. By starting off training inputs having high SNR (≥ 100) and then gradually increasing the noise in each subsequent training session, we found that the performance of prediction can be quickly maximized for low SNR while retaining performance at high SNR. We first trained the predictor on the datasets labeled with the BBH masses and then copied the weights of this network to initialize the classifier and then trained it on datasets having 90% pure random noise inputs, after adding a softmax layer. This transfer learning procedure, similar to multi-task learning, decreases the training time for the classifier and improves its sensitivity.

III. RESULTS

The sensitivity (probability of detecting a true signal) of the classifier as a function of SNR is shown in Fig. 2. We achieved 100% sensitivity when SNR is greater than 10. The false alarm rate was tuned to be less than 1%, i.e., 1 per 100 seconds of noise in our test set was classified as signals. Given independent noise from multiple detectors, this implies our 2-detector false alarm rate would be less than 0.01%, when the classifier is applied independently to each detector and coincidence

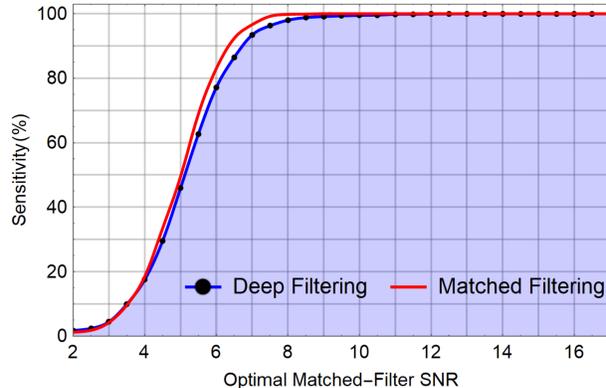


FIG. 2. **Sensitivity of detection with real LIGO noise.** The curve shows the sensitivity of detecting GW signals injected in real LIGO noise from our test set using `Deep Filtering` and with matched filtering with the same template bank used for training. The SNR is on average proportional to 10 ± 1.5 times the ratio of the amplitude of the signal to the standard deviation of the noise. This implies that we can detect signals much weaker than the background noise. While matched-filtering has the advantage of being optimized with the PSD of the LIGO noise in the test set, `Deep Filtering` was only trained on noise from other events, therefore our results demonstrate that CNNs automatically generalize to non-stationary LIGO noise having different PSDs without retraining.

is enforced. Although the false alarm rate can be further decreased by tuning the fraction of noise used for training or by checking that the predicted parameters are consistent, this may not be necessary since running matched-filtering pipelines with a few templates close to our predicted parameters can quickly eliminate these false alarms.

Our predictor was able to successfully measure the component masses given noisy GW signals, that were not used for training, with an error lower than the spacing between templates for optimal matched-filter SNR ≥ 15.0 . The variation in relative error against SNR is shown in Fig. 3. We observed that the errors follow a Gaussian distribution for each region of the parameter space for SNR greater than 10. For high SNR, our predictor achieved mean relative error less than 10%, whereas matched-filtering with the same template bank always has error greater than 10%. This implies that `Deep Filtering` is capable of interpolating between templates seen in the training data.

Although, we trained only on simulated quasi-circular non-spinning GW injections, we applied `Deep Filtering` to the LIGO data streams containing a true GW signal, GW150914, using a sliding window of 1s width with offsets of 0.2s through the data around each event from each detector. This signal was correctly identified by the classifier at the true position in time and each

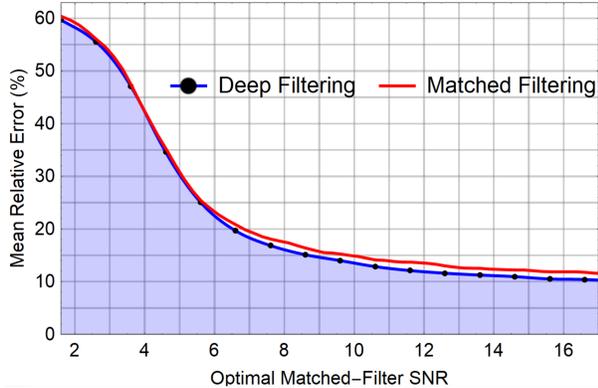


FIG. 3. **Error in parameter estimation with real LIGO noise.** This shows the mean percentage absolute error of estimating masses on our testing signals at each SNR, injected in real LIGO noise from events not used for training, compared to matched filtering using the same template bank that was used for training. While the mean error of matched-filtering, with the same template bank used for training, is always greater than 11% at every SNR we can see that `Deep Filtering` can interpolate to test set signals with intermediate parameter values.

of the predicted component masses were within the published error bars [1]. There were zero false alarms after enforcing the constraint that the detection should be made simultaneously in multiple detectors. This shows that deep learning is able to generalize to real GW signals after being trained only with simulated GW templates injected into LIGO noise from events with different PSDs.

The data from the first LIGO event, that was used for testing, contained a large number of non-Gaussian transient noise called glitches. Therefore, our results demonstrate that the `Deep Filtering` method can automatically recognize these glitches and classify them as noise. This suggests that by adding additional neurons for each “glitch” class, `Deep Filtering` could serve as an alternative to glitch classification algorithms [20, 21, 27].

Furthermore, we conducted some experiments to show the resilience of `Deep Filtering` to transient disturbances, with a simulated set of sine-Gaussian glitches, which cover a broad range of morphologies found in real LIGO glitches [28]. We ensured that a different set of frequencies, amplitudes, peak positions, and widths were used for training and testing. We then injected some of these glitches into the training process and found that the classifier CNN was able to easily distinguish new glitches from true signals, with a false alarm rate less than 1%, using only single detector inputs. When we applied the standard naive matched-filtering algorithm to the same test set of glitches, approximately 30% of glitches were classified as signals due to their high SNR. This is because matched-filtering is unable to distinguish signals from loud glitches having similar frequencies.

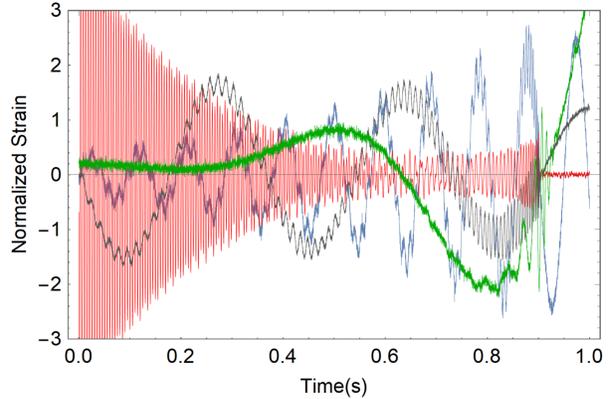


FIG. 4. **Detecting signals contaminated by glitches.** These are some of the signals in our test set injected into real LIGO noise and superimposed with simulated sine-Gaussian glitches from the test set. Each of these inputs were correctly detected as a signal by our classifier. This indicates that `Deep Filtering` can be used as an automatic trigger generator for GW signals that occur in coincidence with glitches.

Note that additional signal consistency tests and coherence across detectors can be enforced to decrease this false alarm rate for both methods.

We then tested the performance of `Deep Filtering`, when a signal happens to occur in coincidence with a glitch, i.e., the signal is superimposed with both a glitch and real LIGO noise (see Fig. 4). We trained the CNNs by injecting glitches from the training set and measured the sensitivity of the classifier on the test set signals superimposed with glitches sampled from the test set of glitches. We found that over 80% of the signals with SNR of 10 were detected, and their parameters estimated with less than 30% relative error, even after they were superimposed with glitches. These results are very promising, motivating an in-depth investigation, since we may be able to automatically detect GW signals that occur during periods of bad data quality in the detectors using `Deep Filtering`, whereas currently such periods are often vetoed and left out of the analysis by LIGO pipelines.

Another important experiment that we carried out was to inject waveforms obtained from simulations of eccentric BBH systems, with eccentricities between 0.1 and 0.2 when entering the LIGO band that we performed using the open-source Einstein Toolkit [24] on Blue Waters, as well as waveforms from spin-precessing binaries from the public SXS catalog [29]. We found that, after injecting these signals into real LIGO noise from the test set events, they were detected by our classifier with the same sensitivity as the original test set of quasi-circular BBH waveforms, with 100% sensitivity of detection and less than 35% error in estimating masses for $\text{SNR} > 10$, thus demonstrating its ability to automati-

cally generalize to GW signals from new classes of BBH mergers beyond the training data to some extent. This is particularly promising since recent studies indicate that moderately eccentric BBH populations may be missed by quasi-circular GW searches. Parameter estimation on these new classes of signals could be more challenging, without training, as spin and eccentricity may have the effect of mimicking features of systems of different masses. We expect that by training on the full parameter space of signals, we will be able to decrease this error.

Both our CNNs are only 23MB in size each, yet encodes all the relevant information from about 2,500 GW templates (~300MB) of templates and several GB of noise used to generate the training data. The time-domain matched-filtering algorithm used for comparison required over 2s to analyze 1s inputs on our CPU. The average time taken for evaluating each of our CNNs per second of data is approximately 85 milliseconds and 540 microseconds using a single CPU core and GPU respectively, thus enabling analysis even faster than real-time. While the computational cost of matched-filtering grows exponentially with the number of parameters, the `Deep Filtering` algorithm requires only a one-time training process, after which the analysis can be performed in constant time. Therefore, we expect the speed-up compared to matched-filtering to further increase by several orders of magnitude when the analysis is extended to larger number of parameters. When considering the full range of signals that span a very high-dimensional parameter space which cannot be sampled densely due to computational costs, we expect `Deep Filtering` may have higher sensitivity due to its ability to interpolate and because the one-time training process can be carried out with template banks much larger than what is feasible to use with matched-filtering.

IV. CONCLUSION

In this article, we have shown for the very first time that CNNs can be used for both *detection* and *parameter estimation* of GW signals in LIGO data. This new paradigm for real-time GW analysis may overcome outstanding challenges regarding the extension of established GW detection algorithms to higher dimensions for targeting a deeper parameter space of astrophysically motivated GW sources. The results of `Deep Filtering` can be quickly verified, and the time of coalescence computed, via matched-filtering with a small set of templates in the region of parameter space predicted by the CNN. Therefore, by combining `Deep Filtering` with well-established GW detection algorithms we may be able to accelerate multimessenger campaigns, pushing the frontiers of GW astrophysics and fully realize its potential for scientific discovery.

The intrinsic scalability of deep learning can overcome the curse of dimensionality [13] and take advantage of massive datasets. This ability could enable fast automated GW searches covering millions or billions of templates over the full range of parameter-space that is beyond the reach of existing algorithms. Extending `Deep Filtering` to predict any number of parameters such as spins, eccentricities, etc., or additional classes of signals or noise, is as simple as adding an additional neuron for each new parameter, or class, to the final layer and training with noisy waveforms with the corresponding labels. Furthermore, the input dimensions of the CNNs can be enlarged to take time-series inputs from multiple detectors, thus allowing coherent searches and measurements of parameters such as sky locations.

The results presented in this article provide a strong incentive to extend `Deep Filtering` to cover the parameter space of spin-aligned BBHs on quasi-circular orbits and beyond. In addition to our primary results, we have also presented experiments exhibiting the remarkable resilience of this method for detection in periods of bad data quality, even when GW signals are contaminated with non-Gaussian transients. This motivates including additional classes of real glitches, e.g., from the Gravity Spy project [21], to the training process to automatically classify and cluster glitches directly from the time-series inputs. It may be possible that by investing more effort in hyperparameter tuning and by experimenting with deeper architectures we may be able to outperform the sensitivity of matched-filtering searches in the future (since matched-filtering is proven to be the optimal algorithm only for Gaussian noise). Therefore, a single, robust, and efficient data analysis pipeline for GW detectors, based on `Deep Filtering`, that unifies detection and parameter estimation along with glitch classification and clustering in real-time with very low computational overhead may be built and deployed in the next observing run of LIGO in the near future.

`Deep Filtering` can be extended to provide instant alerts with accurate parameter estimates for EM follow-up campaigns. It is also possible to apply CNNs to obtain uncertainty distributions for the parameters [30], which will be explored in a following article. As deep CNNs excel at image processing, applying the same approach to analyze raw telescope data may also accelerate the subsequent search for transient EM counterparts. Our results also suggest that, given templates of signals, `Deep Filtering` can be used as a generic tool for efficiently detecting and estimating properties of highly noisy time-domain signals embedded in Gaussian noise or non-stationary non-Gaussian noise, even in the presence of transient disturbances. Therefore, we anticipate that the techniques we have developed for analyzing weak signals in complex noise backgrounds could be useful in other domains of science and technology.

V. ACKNOWLEDGEMENTS

This research is part of the Blue Waters sustained-petascale computing project, supported by the NSF (awards OCI-0725070 and ACI-1238993), NCSA, and the state of Illinois. The eccentric numerical relativity simulations used in this article were generated with the open source, community software, the Einstein Toolkit on Blue Waters and XSEDE (TG-PHY160053). We ex-

press our gratitude to Gabrielle Allen, Ed Seidel, Roland Haas, Miguel Holgado, Haris Markakis, Zhizhen Zhao and other members of the **NCSA Gravity Group** along with Prannoy Mupparaju for comments and interactions. We acknowledge the LIGO/Virgo collaboration, Wolfram Research, and NVIDIA for their support. We thank Vlad Kindratenko for granting us access to the Innovative Systems Lab at NCSA.

-
- [1] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, and et al., *Physical Review Letters* **116**, 061102 (2016), [arXiv:1602.03837 \[gr-qc\]](#).
- [2] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, and et al., *Physical Review Letters* **116**, 241103 (2016), [arXiv:1606.04855 \[gr-qc\]](#).
- [3] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, and et al., *Physical Review X* **6**, 041015 (2016), [arXiv:1606.04856 \[gr-qc\]](#).
- [4] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, et al., *Physical Review Letters* **118**, 221101 (2017).
- [5] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, and et al., *Physical Review Letters* **119**, 161101 (2017), [arXiv:1710.05832 \[gr-qc\]](#).
- [6] D. Eichler, M. Livio, T. Piran, and D. N. Schramm, *Nature* **340**, 126 (1989).
- [7] R. Narayan, B. Paczynski, and T. Piran, *Astrophysical Journal Letters* **395**, L83 (1992), [astro-ph/9204001](#).
- [8] S. A. Usman, A. H. Nitz, I. W. Harry, C. M. Biwer, D. A. Brown, M. Cabero, C. D. Capano, T. Dal Canton, T. Dent, S. Fairhurst, M. S. Kehl, D. Keppel, B. Krishnan, A. Lenon, A. Lundgren, A. B. Nielsen, L. P. Pekowsky, H. P. Pfeiffer, P. R. Saulson, M. West, and J. L. Willis, *Classical and Quantum Gravity* **33**, 215004 (2016), [arXiv:1508.02357 \[gr-qc\]](#).
- [9] E. A. Huerta, P. Kumar, B. Agarwal, D. George, H.-Y. Schive, H. P. Pfeiffer, R. Haas, W. Ren, T. Chu, M. Boyle, D. A. Hemberger, L. E. Kidder, M. A. Scheel, and B. Szilágyi, *Physical Review D* **95**, 024038 (2017).
- [10] L. Lehner and F. Pretorius, *ARA&A* **52**, 661 (2014), [arXiv:1405.4840 \[astro-ph.HE\]](#).
- [11] E. A. Huerta, P. Kumar, S. T. McWilliams, R. O’Shaughnessy, and N. Yunes, *Physical Review D* **90**, 084016 (2014), [arXiv:1408.3406 \[gr-qc\]](#).
- [12] Y. Lecun, Y. Bengio, and G. Hinton, *Nature* **521**, 436 (2015).
- [13] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016).
- [14] J. Schmidhuber, *Neural Networks* **61**, 85 (2015).
- [15] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al., *Nature* **550**, 354 (2017).
- [16] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, *ArXiv e-prints* (2016), [arXiv:1609.03499 \[cs.SD\]](#).
- [17] Y. D. Hezaveh, L. P. Levasseur, and P. J. Marshall, *Nature* **548**, 555 (2017).
- [18] N. Sedaghat and A. Mahabal, *ArXiv e-prints* (2017), [arXiv:1710.01422 \[astro-ph.IM\]](#).
- [19] K. A. Pearson, L. Palafox, and C. A. Griffith, *ArXiv e-prints* (2017), [arXiv:1706.04319 \[astro-ph.IM\]](#).
- [20] D. George, H. Shen, and E. A. Huerta, *ArXiv e-prints* (2017), [arXiv:1706.07446 \[gr-qc\]](#).
- [21] M. Zevin, S. Coughlin, S. Bahaadini, E. Besler, N. Rohani, S. Allen, M. Cabero, K. Crowston, A. Katsaggeolos, S. Larson, T. K. Lee, C. Lintott, T. Littenberg, A. Lundgren, C. Oesterlund, J. Smith, L. Trouille, and V. Kalogera, *ArXiv e-prints* (2016), [arXiv:1611.04596 \[gr-qc\]](#).
- [22] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, *Proceedings of the IEEE* **86**, 2278 (1998).
- [23] A. Taracchini, A. Buonanno, Y. Pan, T. Hinderer, M. Boyle, D. A. Hemberger, L. E. Kidder, G. Lovelace, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel, B. Szilágyi, N. W. Taylor, and A. Zenginoglu, *Physical Review D* **89**, 061502 (2014), [arXiv:1311.2544 \[gr-qc\]](#).
- [24] F. Löffler, J. Faber, E. Bentivegna, T. Bode, P. Diener, R. Haas, I. Hinder, B. C. Mundim, C. D. Ott, E. Schnetter, G. Allen, M. Campanelli, and P. Laguna, *Classical and Quantum Gravity* **29**, 115001 (2012), [arXiv:1111.3344 \[gr-qc\]](#).
- [25] F. Yu and V. Koltun, in *ICLR* (2016).
- [26] D. P. Kingma and J. Ba, *CoRR abs/1412.6980* (2014).
- [27] N. Mukund, S. Abraham, S. Kandhasamy, S. Mitra, and N. S. Philip, *Phys. Rev. D* **95**, 104059 (2017).
- [28] J. Powell, D. Trifirò, E. Cuoco, I. S. Heng, and M. Cavaglià, *Classical and Quantum Gravity* **32**, 215012 (2015), [arXiv:1505.01299 \[astro-ph.IM\]](#).
- [29] T. Chu, H. Fong, P. Kumar, H. P. Pfeiffer, M. Boyle, D. A. Hemberger, L. E. Kidder, M. A. Scheel, and B. Szilágyi, *Classical and Quantum Gravity* **33**, 165001 (2016), [arXiv:1512.06800 \[gr-qc\]](#).
- [30] L. Perreault Levasseur, Y. D. Hezaveh, and R. H. Wechsler, *Astrophysical Journal Letters* **850**, L7 (2017).