

Physics-inspired deep learning to characterize the signal manifold of quasi-circular, spinning, non-precessing binary black hole mergers

Asad Khan^{a,b,*}, E.A. Huerta^{a,b,c}, Arnav Das^{a,d}

^a National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

^b Department of Physics, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

^c Department of Astronomy, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

^d Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

ARTICLE INFO

Article history:

Received 20 April 2020

Received in revised form 11 June 2020

Accepted 14 July 2020

Available online 27 July 2020

Editor: H. Peiris

Keywords:

Physics-inspired AI

Black hole mergers

Gravitational waves

ABSTRACT

The spin distribution of binary black hole mergers contains key information concerning the formation channels of these objects, and the astrophysical environments where they form, evolve and coalesce. To quantify the suitability of deep learning to estimate the individual spins, effective spin and mass-ratio of quasi-circular, spinning, non-precessing binary black hole mergers, we introduce a modified version of WaveNet trained with a novel optimization scheme that incorporates general relativistic constraints of the spin properties of astrophysical black holes. The neural network model is trained, validated and tested with 1.5 million $\ell = |m| = 2$ waveforms generated within the regime of validity of NRHybSur3dq8, i.e., mass-ratios $q \leq 8$ and individual black hole spins $|s|_{(1,2)} \leq 0.8$. To reduce time-to-insight, we deployed a distributed training algorithm at the IBM Power9 Hardware-Accelerated Learning cluster at the National Center for Supercomputing Applications to reduce the training stage from 1 month, using a single V100 NVIDIA GPU, to 12.4 hours using 64 V100 NVIDIA GPUs. We have also fully trained this model using 1536 V100 GPUs (256 nodes) in the Summit supercomputer at Oak Ridge National Laboratory, achieving state-of-the-art accuracy within just 1.2 hours. Using this neural network model, we quantify how accurately we can infer the astrophysical parameters of black hole mergers in the absence of noise. We do this by computing the overlap between waveforms in the testing data set and the corresponding signals whose mass-ratio and individual spins are predicted by our neural network. We find that the convergence of high performance computing and physics-inspired optimization algorithms enable an accurate reconstruction of the mass-ratio and individual spins of binary black hole mergers across the parameter space under consideration. This is a significant step towards an informed utilization of physics-inspired deep learning models to reconstruct the spin distribution of binary black hole mergers in realistic detection scenarios.

© 2020 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>). Funded by SCOAP³.

1. Introduction

Gravitational wave (GW) observations provide unique insights into the formation channels of compact binary systems. For instance, it is expected that inferring orbital eccentricity in GW observations [1–13] may provide the most conclusive evidence for the existence of compact binary systems in dense stellar environments [14–34]. For instance, in the case of binary black hole (BBH) mergers, it is assumed that the spin distribution of BBHs formed in dense stellar environments may be distributed isotropically, whereas BBHs formed through massive stellar evolution in

isolation may have spin distributions that are aligned with the binary's orbital angular momentum [35,36].

As the number of GW observations of BBH mergers continues to grow in years to come [37–39], it will be possible to infer the astrophysical properties of these sources and elucidate their formation history.

The goal of this article is to explore how deep learning handles parameter space degeneracies in the signal manifold of quasi-circular, spinning, non-precessing BBH mergers; and to quantify how accurately deep learning may constrain the individual spins, effective spin and mass-ratio of these GW sources in the absence of noise. We then go on to discuss the computational grand challenges that naturally arise when one tries to address these problems: (1) the parameter space that needs to be sampled is very

* Corresponding author.

E-mail address: khan74@illinois.edu (A. Khan).

large, requiring the use of TB-size waveform datasets, and thereby demanding the development of novel distributed algorithms to use many GPUs to fully train deep learning algorithms in a reasonable amount of time; and (2) the need to incorporate domain knowledge into the optimization of deep learning algorithms to accelerate their convergence and to ensure that their predictions are physically consistent.

In connection to the first challenge mentioned above, we introduce distributed training algorithms that reduces the training stage of the neural network model used in this study from one month using a single V100 GPU to: (i) 12.4 hours using 64 V100 GPUs at the Hardware Accelerated Learning (HAL) cluster at the National Center for Supercomputing Applications (NCSA); and (ii) within 1.2 hours using 1536 V100 GPUs at the Summit supercomputer at Oak Ridge National Laboratory. These results establish a record in the number of GPUs used to train these types of physics-inspired deep learning models.

Regarding the second challenge, we have found that naive methods to train deep learning architectures lead to rather sub-optimal results. However, we show that when we use physics-inspired optimization algorithms, which incorporate general relativistic constraints of the spin of BBHs, we are able to accurately recover the individual spins and mass-ratio of BBH signals across the mass-ratio under consideration. This analysis provides benchmarks for the performance of deep learning models when applied to the reconstruction of these parameters in the absence of noise, and provides a baseline of accuracy when real noise from GW observations is taken into consideration. That study will be presented shortly in a follow up paper.

To describe the signal manifold of quasi-circular, spinning, non-precessing BBH mergers, we use a catalog of over one million time-series waveforms that are parameterized in natural units, i.e., we use numerical relativity (NR) type waveforms, such that the GWs that describe BBH mergers may be fully described in terms of mass-ratio, q , and the individual spins of the binary components, s_i^z with $i = \{1, 2\}$. The neural network will be trained to cover this 3-D signal manifold (q, s_1^z, s_2^z) , with the aim of accurately inferring these three parameters when unlabelled waveforms are fed into the neural network.

This work builds upon an emergent field of deep learning research which has thus far provided new methodologies to do detection and point-parameter estimation for GW sources in the context of simulated noise [40], and real advanced LIGO noise [41,42]; detection-only methods in the context of simulated noise [43,44], real noise [45]; denoising of GW signals [46–48], among others (see [49] and references therein). In the context of other recent deep learning studies

- Our work differs from [50] in that they focus on the physics of waveform signals during the ringdown phase, whereas we are interested in the effective spin or individual spins of BBH mergers through the waveform evolution.
- The authors in [51] use conditional variational autoencoders to produce fast posterior probability estimates for *non-spinning* BBH mergers embedded in simulated noise. Furthermore, the authors in [52] introduce the use of auto-regressive normalizing flows for rapid inference. They use a small data set of spinning BBH mergers, and assume signals with signal-to-noise ratios $\gtrsim 20$, to produce marginalized posterior distributions for a number of parameters, including the individual spins and effective spin of non-precessing BBH mergers in the context of simulated noise. In contrast, in this paper we focus on the development of physics-inspired neural network models that, in the absence of noise, provide insights about how deep learning handles parameter space degeneracies and the measurement of individual spins, effective spin and mass-ratio of

spinning, non-precessing BBH mergers. In brief, herein we develop a computational framework to minimize time-to-insight by combining extreme scale computing and physics-inspired models. This approach will be exploited in a future study to carefully assess how simulated and real noise bias the results we present in this article.

The previous list shows that it is timely and relevant to start developing deep learning models that encapsulate as much physics as possible, which may only be accomplished by using larger training datasets and developing physics-inspired models and optimization schemes. In turn, this demands the development of distributed training algorithms in high performance computing platforms to reduce time-to-solution. This article is a significant step in that direction.

This study is organized as follows. Section 2 describes the data sets used to train, validate and test our neural network model. We also describe therein the deep learning model used, and show how to build a physics-inspired optimization algorithm that significantly improves the predictive accuracy of our neural network. Our findings are presented and discussed in Section 3. We summarize this work and outline future directions of work in Section 4.

2. Methods

In this section we describe the data sets used to train, validate and test our neural network model. We also describe the neural network architecture used for this work, and the construction of a physics-inspired optimization algorithm that significantly improves the ability of our neural network to accurately infer the astrophysical properties of BBH mergers from the time-series NR waveforms that describe them.

2.1. Data curation

We generate our training, validation and test sets using NRHyb-Sur3dq8 [53], a surrogate model for hybridized non-precessing NR waveforms. While this surrogate model may be extrapolated up to mass-ratio $q = 10$ and $s_i^z = 0.998$, with $i = \{1, 2\}$, it has only been trained with 104 NR waveforms in the parameter range $q \leq 8$ and $s_i^z \leq 0.8$, and hence we restrict our data sets to the same parameter span. Furthermore, we consider the $\ell = |m| = 2$ mode to train, validate and test our neural network model. To be consistent with the data used to train, validate and test our neural network, throughout this paper we use a geometric unit system in which $G = c = 1$.

The signals produced for this analysis are such that the amplitude peak occurs at $t = 0M$, covering a time span $t \in [-10, 000M, 130M]$ with a time step $\Delta t = 0.1M$. A sample waveform is shown in Fig. 1.

The training set is generated by sampling the mass-ratio $q \in [1, 8]$ in steps of $\Delta q = 0.08$; and the individual spins $s_i^z \in [-0.8, 0.8]$ in steps of $\Delta s_i^z = 0.012$. This is equivalent to ~ 1.5 million waveforms. The validation and test sets are generated by alternately sampling the intermediate values, i.e. by sampling q and s_i^z in steps of 0.16 and 0.024 to lie between training set values, for a total of $\sim 190,000$ waveforms each respectively. The distributions of parameters for training, validation and test sets is summarized in Fig. 2. The entire data set is ~ 1.5 TB in size, and we make use of `mpi4py` [54–56] to parallelize the data generation using the Campus Cluster at the University of Illinois at Urbana-Champaign [57].

2.2. Neural network model architecture and loss function

The neural network architecture consists of two fundamental components, a shared base/root consisting of layers slightly mod-




Fig. 1. Sample waveform for a binary black hole system with mass-ratio $q = 7.9$, and whose binary components have spins $s_1^z = s_2^z = 0.8$, respectively. All waveforms used in this analysis cover the time range $t \in [-10, 000 \text{ M}, 130 \text{ M}]$, and are sampled with a time step $\Delta t = 0.1 \text{ M}$.




Fig. 2. Sampling of the signal manifold $q \in [1, 8]$, $s_{[1,2]}^z \in [-0.8, 0.8]$ to construct the training (light blue dots), validation (dark blue dots) and testing (red dots) data sets.

ified from the WaveNet [58] architecture, and two branches consisting of fully connected layers that take in features extracted from the root to predict the mass ratio and the individual spins of the binary components, respectively.

WaveNet is a probabilistic and auto-regressive model originally released by DeepMind for generating raw audio waveforms, with predictive distribution for each audio sample conditioned on all previous ones. Trained on data with tens of thousands of samples per second of audio, it exhibited not only state of the art performance on text-to-speech, but also showed promising results in music generation and as a discriminative model for phoneme/speech recognition. Additionally WaveNet inspired architectures have also been demonstrated to successfully denoise recent GW observations [47]. Inspired by such recent successes, as well as key architectural features (which we describe in more detail below) suited to processing wideband raw waveforms, we modify the WaveNet architecture and test it as a discriminative model to predict key parameters of GWs from quasi-circular, spinning, non-precessing binary black hole mergers.

The key architectural components of a WaveNet are dilated causal convolutions, gated activation units, and the usage of residual and skip connections, which we describe below.




Fig. 3. Stack of dilated convolutional layers. Zero padding on either side is not shown.

2.2.1. Root layers

Dilated causal convolutions. The original WaveNet was used as a generative model to predict audio sample x_t conditioned on all previous timesteps, where the joint probability of the waveform $\mathbf{x} = \{x_1, x_2, \dots, x_T\}$ is factorized as a product of conditional probabilities $p(\mathbf{x}) = \prod_{t=1}^T p(x_t | x_1, \dots, x_{t-1})$. Traditionally recurrent neural networks (RNNs) and Long short-term memory (LSTM) models have been used to model conditional probabilities of sequences, but they are typically considerably slower than convolutional networks and also suffer from vanishing gradient problem for very long sequences. WaveNet addressed the aforementioned issues using causal convolutions, which are implemented by shifting the output of a normal convolution by a few timesteps, and ensuring that the prediction made by the model at time step t does not depend on the future timesteps $x_{t+1}, \dots, x_T$. Since ours is a discriminative model, we turn off the causality of the WaveNet architecture, so that our model can simultaneously process all time steps of the waveform.

A dilated convolution is an effective way to increase the receptive field of a convolutional network by orders of magnitude with only a small computational overhead. This is achieved by applying a convolutional filter over the input sequence by skipping input values with a certain step size (called the ‘dilation’). Fig. 3 shows dilated convolutions for dilations of 1, 2, 4, and 8 respectively. Similar to the original WaveNet architecture, we double the dilation for every layer up to a limit and then repeat the same stack of layers, to efficiently capture long range structure of the waveform.

Gated activation units. A potential advantage of using LSTMs is that they have multiplicative units (memory gates) that may help with modeling complex inter-dependencies between time steps. In the original WaveNet architecture, this is amended by replacing the rectified linear units (ReLU) between convolutions with the following gated activation unit:

$$\mathbf{z} = \tanh(W_{f,k} * \mathbf{x}) \odot \sigma(W_{g,k} * \mathbf{x}), \quad (1)$$

where $*$ is a convolution operator, $\odot$ is an element-wise multiplication operator, $\sigma(\cdot)$ is the sigmoid function, k is the layer index, and W_f and W_g denote filter and gate convolutions, respectively. We keep the same gated activation Units in our network.

Residual and skip connections. Similar to the original WaveNet architecture, both residual and skip connections are used for fast convergence by resolving vanishing gradients when training deep neural networks. Fig. 4 shows a residual block of the model, which is stacked many times in the network. More in depth discussion of the structure of WaveNet may be found in the original paper [59].

2.2.2. Leaf layers

At the end of root layer the model outputs a sequence with the same time dimensionality as the input sequence. Since the input




Fig. 4. Network architecture. Residual blocks of Wavenet (left) and leaf layers (right).

sequence has 101,300 time-steps, we pick the last 10,000 time-steps from the output sequence, flatten and feed into two leafs of fully connected layers predicting the mass ratio q and spins s_i^z respectively, as shown in Fig. 4. Hence the shared root can be thought of as a feature extractor, and the leafs as sub-networks specializing to predict different parameters from the extracted features. The motivation behind this root/leaf structure is two-fold; firstly the mass ratio $q \in [1, 8]$ and the spins $s_i^z \in [-0.8, 0.8]$ have different numerical ranges, and secondly we performed experiments without a leaf structure and empirically got sub-optimal results.

2.2.3. Loss function

We performed several experiments to find an optimal way to constrain s_i^z predictions from the second leaf sub-network. In the first iteration, we predicted s_i^z directly using mean-squared error as the loss function. This resulted in reasonable predictions for s_1^z , but s_2^z predictions could not be constrained with continued training. Consequently, we explored the use of effective one-body general relativistic dynamics of two spinning black holes as delineated in [60]. Specifically, we focus on the derivation of a spin-dependent effective one-body Hamiltonian for small and moderate spins as a ν -deformation of a Kerr metric of mass $M \equiv m_1 + m_2$ and effective spin

$$S_{\text{eff}} = \sigma_1 s_1^z + \sigma_2 s_2^z, \quad (2)$$

where $\sigma_1 \equiv 1 + \frac{3}{4q}$ and $\sigma_2 \equiv 1 + \frac{3q}{4}$. In addition, we also consider the effective spin parameter [35,61]

$$\sigma_{\text{eff}} = \frac{m_1 s_1^z + m_2 s_2^z}{m_1 + m_2} = \frac{q s_1^z + s_2^z}{1 + q}. \quad (3)$$

Predicting S_{eff} and σ_{eff} , and using our prediction for q to solve Equations (2) and (3), we were able to tightly constrain s_i^z predictions.

3. Results

This section presents our main findings in the following format: we first discuss specific challenges to be addressed, namely, the parameter space degeneracy of the signal manifold under consideration and the need to use a large training data set to densely sample this 3-D signal manifold. We then present qualitative results of the performance of our neural network model, and finalize with a quantitative set of results that provide a thorough description of the realm of applicability of our neural network model.

3.1. Parameter space degeneracy and convergence of high performance computing with deep learning

Inferring key parameters from the signal manifold of quasi-circular, spinning, non-precessing, BBH mergers presents a number of complications given that different time-series NR waveforms, say $h(t)$ and $s(t)$, that describe different astrophysical systems are remarkably similar. To illustrate this property we have selected a few astrophysical systems, and then computed the overlap, $\mathcal{O}(h, s)$, between them and 2-D slices of the signal manifold, determined by the BBH mass-ratio and $s_{\{1,2\}}^z$, as shown in Fig. 5, using the relation

$$\mathcal{O}(h, s) = \max_{t_c, \phi_c} \left(\hat{h} | \hat{s}_{[t_c, \phi_c]} \right), \quad \text{with } \hat{h} = h(h|h)^{-1/2}, \quad (4)$$

where $\hat{s}_{[t_c, \phi_c]}$ indicates that the normalized waveform $\hat{s}$ has been time- and phase-shifted. Using this metric, Fig. 5 shows that the overlap between a given signal and other BBH systems with remarkably different spin combinations are effectively indistinguishable. Other observations we obtain from these results is that inferring the individual spins of near-equal mass-ratio systems is very hard given the intrinsic symmetry of the system, i.e., the binary components may be interchanged $m_1 \iff m_2$, as shown in the top panels of Fig. 5. Indeed, near equal BBH systems present the largest regions of degeneracy across the signal manifold. We also observe that systems with small to moderate spins present large regions of degeneracy, as shown in the second and third row of panels in Fig. 5. These features are also present in BBHs whose binary components have large spin values, as shown in the bottom panels of Fig. 5. Following [62], we recast the results in Fig. 5 in terms of the symmetric mass-ratio, η , and the effective spin, σ_{eff} , where

$$\eta = \frac{m_1 m_2}{(m_1 + m_2)^2} \quad (5)$$

These results provide a holistic perspective on spin and mass-ratio degeneracy across the parameter space under consideration as shown in Fig. 6. In summary, inferring the individual spins of BBH mergers is a rather challenging endeavor.

To try to address these challenges, we begin by densely sampling the signal manifold under consideration following the methodology described in Section 2.1, which leads to the construction of a data set that includes over 1.5×10^6 waveforms. Training the model described in Section 2.2 with this large data




Fig. 5. From top to bottom, each row presents regions of parameter space degeneracy for a random sample of binary black hole mergers with parameters: (i) $q = 1.0$, $s_1^z = 0.8$, $s_2^z = -0.8$; (ii) $q = 6.00$, $s_1^z = 0.40$, $s_2^z = -0.40$; (iii) $q = 7.12$, $s_1^z = 0.70$, $s_2^z = -0.30$; and (iv) $q = 7.92$, $s_1^z = -0.70$, $s_2^z = -0.70$. The color bar indicates the overlap between each of the four aforementioned binary black hole mergers and the three corresponding mass-ratio slices shown on each row.

set using a single V100 GPU would take about 30 days to ensure that the model converges and attains optimal performance, which is quantified by computing the overlap between ground-truth signals and signals whose parameters are predicted by the neural network. In order to minimize time-to-insight and to test multiple, physics-inspired optimization algorithms, we had to design and deploy a data-parallel distributed training scheme on the Hardware-Accelerated Learning (HAL) deep learning cluster at the National Center for Supercomputing Applications [63,64]. This cluster has 64 NVIDIA V100 GPUs distributed evenly across 16 nodes, and connected by NVLink 2.0 inside the nodes and EDR InfiniBand across the nodes. As shown in the left panel of Fig. 7, this approach reduces the training stage to only 12.4 hours. Additionally, we also scaled the data-parallel distributed training strategy

up to 6144 NVIDIA V100 GPUs at 80% efficiency on the Summit supercomputer at Oak Ridge National Laboratory, as shown in right panel of Fig. 7. In data-parallel distribution scheme, the neural network model is replicated on each individual GPU and each replication is fed non-overlapping batches of training data in parallel. After each batch, the GPUs communicate to synchronize gradients and update model weights. Because this scheme involves a linear increase in global batch size with the number of GPUs, it has been observed that such scaling leads to a degradation in convergence and generalization of the model [65]. To address this issue, we employed the layer-wise adaptive large batch optimization technique (LAMB) [66], and successfully trained the model without degradation in convergence using 1536 NVIDIA V100 GPUs within 1.2 hours. We have made the trained model [67] and test-




Fig. 6. As Fig. 5, but now presenting results in terms of symmetric mass ratio, η , and effective spin, σ_{eff} . True parameters are indicated by red stars. The color bar indicates the waveform overlap between true parameters (red star) and the rest of the parameter space.




Fig. 7. Left: Speed up in training obtained by designing and deploying a distributed training scheme at the Hardware-Accelerated Learning (HAL) deep learning cluster at the National Center for Supercomputing Applications. This approach reduces the training stage from one month (using a single V100 GPU) to 12.4 hours by scaling the training to the entire cluster, which consists of 64 NVIDIA V100 GPUs distributed evenly across 16 nodes, and connected by NVLink 2.0 inside the nodes and EDR InfiniBand across the nodes [63]. Right: Speed up obtained upon deploying and tuning our distributed training scheme on the Summit supercomputer at Oak Ridge National Laboratory.

ing dataset publicly available at Data and Learning Hub for Science (DLHub) [68,69] hosted at Argonne National Laboratory.

3.2. Qualitative analysis

The first assessment we have conducted to examine the performance of our fully trained model is presented in Fig. 8. Therein we show the absolute errors in the estimation of the mass-ratio and individual spins for several fixed, mass-ratio slices of the signal manifold, where the absolute error ΔQ_i for any quantity Q_i is defined as:

$$\Delta Q_i = Q_i^{\text{prediction}} - Q_i^{\text{true}} \quad (6)$$

The top left panel shows again the degeneracy expected for near equal mass-ratio systems. Note that while we can obtain a good

estimate of the mass-ratio of these systems, it is hard to infer the individual spins of the binary components. Once this degeneracy is broken, i.e., we consider $q > 1$ BBH mergers, it is now possible to infer the individual spins and the mass-ratio of the BBH mergers with the accuracy shown in the top-right and bottom-left panels in Fig. 8. As expected, as the mass-ratio of the BBH increases, it becomes increasingly difficult to infer the spin of the secondary—see bottom right panel of Fig. 8. This is expected, since for asymmetric mass-ratio BBH mergers, the effect of the secondary in the dynamics of the binary system becomes less dominant, i.e., for a fixed mass-ratio and spin of the primary, it is difficult to tell apart signals when we vary the spin of the secondary. We provide a summary of these results in Table 1.

The last metric included in Table 1 corresponds to the overlap, $\mathcal{O}(h, s)$, between every waveform in the testing dataset, $h(\theta^i)$ with




Fig. 8. Top and middle panels: absolute prediction errors in the recovery of (q , s_1^z , s_2^z) for a sample of mass ratios. Bottom panel: absolute prediction errors in the recovery of symmetric mass-ratio, effective spin and Hamiltonian effective-spin, (η , σ_{eff} , S_{eff}), across the parameter space under consideration.

Table 1

Summary of the distribution of absolute errors in the estimation of mass-ratio and individual spins over the entire parameter space under consideration: $q \in [1, 8]$, $s_i \in [-0.8, 0.8]$. We present minimum, median, maximum, first quartile, (Q1), and third quartile, (Q3), of the error distributions. Note that in addition to (q , s_i^z), we also provide error summaries for the effective spin, σ_{eff} , Hamiltonian effective-spin, S_{eff} , and overlap between ground-truth signals and signals whose parameters are predicted by our network, $\mathcal{O}$.

	Minimum	Q1	Median	Q3	Maximum
Δq	-0.074	-0.010	-0.004	0.002	0.188
Δs_1	-0.687	-0.006	0.002	0.009	0.441
Δs_2	-0.466	-0.010	-0.000	0.009	0.767
$\Delta \sigma_{\text{eff}}$	-0.074	-0.004	0.002	0.008	0.098
ΔS_{eff}	-0.516	-0.037	-0.004	0.032	0.493
$\mathcal{O}$	0.278	0.986	0.995	0.999	1.000

ground-truth parameters $\theta^i \rightarrow (q, s_i^z)$, and the signal, s , that best describes h according to our neural network model, i.e., $s(\hat{\theta}^i)$ with $\hat{\theta}^i \rightarrow (q^{\text{predicted}}, s_i^{z, \text{predicted}})$. It is necessary to consider this additional metric because absolute errors provide partial information about the accuracy of our neural network to infer the individual spins and mass-ratio of BBH mergers. One can only ascertain whether these predictions are reliable once we directly compare the signals described by the predicted parameters of our neural network model, and those that it is aiming to characterize. This is the theme of the following section.

3.3. Quantitative analysis

While the preceding qualitative analysis suggests that our neural network may be correctly inferring parameters for the signal manifold under consideration, we have also conducted a quantitative approach, i.e., we collected the predictions of the neural network for the mass-ratio and individual spins for each waveform in the testing data set. Thereafter, we generated waveforms with these *predicted* parameters using NRHybSur3dq8. Finally, we computed the overlap between the ground-truth signals in the testing data set, and the waveforms whose parameters are predicted by our neural network. We summarize the results of this analysis for a sample of mass-ratio slices in Fig. 9. These results indicate that we are able to accurately infer the mass-ratio and individual spins over a broad range of the parameter space under consideration. Indeed, both the median and mean overlap results are above $\mathcal{O} \geq 0.99$. They only drop below 0.98 for BBH mergers at the edge of the signal manifold with $q \sim 8$.

We provide a more detailed analysis of the overlap results in Fig. 10, where we show overlap results at every single point of the testing data set for a number of mass-ratio slices. These results show that our neural network model can predict the mass-ratio and individual spins of BBH mergers with excellent accuracy. Close to the edge of the parameter space under consideration, i.e., $q \rightarrow 8$, the neural network has a gradual decrease in accuracy. The bottom panel in Fig. 10 also presents results for the accuracy with which




Fig. 9. Histograms: waveform matches for a sample of test mass ratios.

our neural network model is able to reconstruct the effective spin, σ_{eff} , and symmetric mass-ratio, η , of BBH mergers across the parameter space under consideration.

Fig. 11 presents a visual representation of high, median and low overlaps samples, i.e., 2 sigma below the median overlap. These random samples from the testing data set show that: (i) our neural network can identify signals that reproduce with excellent accuracy the dynamics of near-equal BBH mergers. However, given the parameter space degeneracy of the signal manifold for $q \sim 1$ systems, it is difficult to accurately recover the individual spins of these systems—see left column in Fig. 11; (ii) for BBH systems with $q \neq 1$ —middle column in Fig. 11—we notice that our network can recover with excellent accuracy both the mass-ratio and individual spins; and (iii) systems with asymmetric mass-ratios—right column in Fig. 11—can be characterized with different levels of accuracy. We notice that even for systems whose overlap is two sigma below the median overlap, the neural network is able to tightly constrain the estimate of the individual spins of the systems, as already summarized in Table 1.

The quantitative analysis we have conducted indicates that deep learning is adequate to reconstruct the mass-ratio and individual spins of NR-type time-series signals that span the 3-D signal manifold $q \in [1, 8]$ and $s_{[1,2]}^z \in [-0.8, 0.8]$. Our findings are in good accord with the expectation that it is hard to estimate the individual spins of BBH mergers whose mass-ratios are near the edges of the signal manifold, i.e., $q \sim 1$ and $q \sim 8$. This is expected given the symmetry of the problem at hand for $q \sim 1$, and the fact that the spin of the secondary is difficult to reconstruct for asymmetric mass-ratio BBH mergers. All these results may be interactively perused at [70].

We provide a careful follow up of BBH systems that are not accurately reconstructed by our neural network in Appendix A, which tend to be concentrated at the edges of the signal manifold.

While we have found that our deep learning model performs very well to characterize $q \sim 1$ BBH mergers, as shown in the top left panel of Fig. 9, we leave to future work the construction of deep learning models that include higher-order modes to explore whether spin reconstruction may be improved for the most asymmetric mass-ratio BBH mergers considered in this study.

4. Conclusion

Inferring the spin distribution of BBH mergers will unveil new and detailed information about the formation channels of these objects. It is then timely and relevant to design signal-processing algorithms that are scalable, and which may readily handle a dense sampling of this higher-dimensional signal manifold.

To contribute to this effort, in this paper we introduced a deep learning model to estimate the individual spins, effective spin and mass-ratio of quasi-circular, spinning, non-precessing, BBH mergers. To do this, we densely sampled the (q, s_1^z, s_2^z) parameter space with over 1.5M waveforms produced with the NRHybSur3dq8 model. We then designed and deployed a distributed training algorithm to reduce the training stage from one month, using a single V100 GPU, to just 12.4 hrs using 64 V100 GPUs in the HAL cluster at the University of Illinois. The convergence of deep learning and high performance computing proved critical to reduce time-to-insight, and enabled us to test a variety of optimization algorithms that incorporate general relativistic constraints. This approach enabled us to demonstrate that, in the absence of noise, physics-inspired deep learning models can effectively reconstruct the mass-ratio and individual spins of the binary components of BBH systems.

Our study sheds light on the ability of deep learning to characterize waveform signals that span a degenerate signal manifold. For instance, in the case of equal or comparable mass-ratio systems we showed that our neural network can reconstruct the mass-ratio of these systems with excellent accuracy. However, the intrinsic symmetry of these systems, i.e., the binary components are indistinguishable $m_1 \leftrightarrow m_2$, reduces the ability of the neural network to provide good estimates of the individual spins of the system. Despite these difficulties, we have found that the signals predicted by our network reproduce the dynamics of the ground-truth signals with excellent accuracy, i.e. we observe a median overlap of 99.6% between the predicted and the ground-truth signals for near equal mass ratios. On the other hand, the neural network can tightly constrain both the mass-ratio and individual spins for asymmetric mass-ratio systems. Near the edge of the signal manifold, i.e., $q \sim 8$ BBH mergers, the neural network has




Fig. 10. Each point in the top and middle panels represents the overlap between a signal in the testing data set and its counterpart whose individual spins and mass-ratio are predicted by our neural network model. The mass-ratio slices presented in this figure were randomly selected from the testing data set. The bottom panel summarizes the accuracy of our neural network model to infer the effective spin and symmetric mass-ratio, (σ_{eff} , η), across the parameter space under consideration.

a drop in accuracy recovering the spin of the secondary. This is expected because the spin of the secondary has a less dominant effect in the overall dynamics of the waveform signal. We leave to future work a systematic assessment of the importance of including higher-order modes to better constrain the spin of the secondary.

In summary, these results represent a significant step towards the design of scalable, physics-inspired neural network models that may be used to infer the spin distribution of actual GW observations. This analysis will enable a firm understanding on the impact of real GW noise when deep learning is used to infer the spin distribution of BBH mergers. An analysis of this nature will be presented in a forthcoming publication.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We thank Arjun Shankar, Tom Gibbs, Junqi Yin, and Jeff Larkin for their support and guidance using the Summit supercomputer.

We also thank Ben Blaiszik, Ryan Chard and Logan Ward for their support deploying our neural network model and testing dataset at the Deep Learning Hub hosted by Argonne National Lab.

We thank Elise Jennings for her guidance and support using the Cooley and Theta supercomputers at Argonne National Lab.

AK and EAH gratefully acknowledge National Science Foundation (NSF) awards OAC-1931561 and OAC-1934757. We are grateful to NVIDIA for donating several Tesla P100 and V100 GPUs that we used for our analysis. AD acknowledges support from the National Center for Supercomputing Applications (NCSA) Students Pushing INnovation (SPIN) undergraduate internship program. The authors thank Dawei Mu and Chit Khin for their superb support using the HAL cluster and the ICCP, respectively.

This work utilized resources supported by the NSF's Major Research Instrumentation program, the Hardware-Learning Accelerated (HAL) cluster, grant OAC-1725729, as well as the University of Illinois at Urbana-Champaign. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

This work made use of the Illinois Campus Cluster, a computing resource that is operated by the Illinois Campus Cluster Program (ICCP) in conjunction with the NCSA and which is supported by funds from the University of Illinois at Urbana-Champaign. This re-




Fig. 11. Random sample of high (top panel), median (middle panel) and low (bottom panel) overlap $\mathcal{O}$ predictions for $q = 1.04$ (left column), 4.08 (middle column), 7.92 (right column). Low overlaps are two sigma below mean overlap for each q .




Fig. A.12. Left: all test set points with Overlap < 0.85 . Right: overlap distribution for the entire test set.

search used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. We also acknowledge NSF grant TG-PHY160053, which provided us access to the Extreme Science and Engineering Discovery Environment (XSEDE) Bridges-AI system. We thank the [NCSA Gravity Group](#) for useful feedback.

Appendix A. Overlap distribution

We have conducted a detailed analysis of the regions of parameter space where our neural network model does not perform optimally. By the numbers, the test data set has 197,516 waveforms. Our neural network predicts signals whose overlap with the corresponding ground-truth signals in the test set has the following distribution:

- Test waveforms with $\mathcal{O} \leq 0.97$: 18,315
- Test waveforms with $\mathcal{O} \leq 0.90$: 817
- Test waveforms with $\mathcal{O} \leq 0.85$: 181

These results are presented at a glance in Fig. A.12. We have also found that if we excise the edges of the signal manifold, i.e., we consider the ranges $q \in [1, 8]$, $s_{1,2}^z \in (-0.8, 0.8)$, then most of the low overlaps are removed.

The reader may conduct a more detailed inspection of these results at [70]. Furthermore, as mentioned in the main body of this article, it is worth exploring whether the inclusion of higher-order modes may improve the performance of this model for asymmetric mass-ratio BBH mergers. Such a study will be pursued in the near future.

References

- [1] F. Antonini, N. Murray, S. Mikkola, *Astrophys. J.* 781 (2014) 45, arXiv:1308.3674 [astro-ph.HE].
- [2] E.A. Huerta, J.R. Gair, *Phys. Rev. D* 79 (2009) 084021, arXiv:0812.4208.
- [3] F. Antonini, M. Gieles, A. Gualandris, arXiv e-prints, arXiv:1811.03640, 2018, arXiv:1811.03640 [astro-ph.HE].
- [4] F. Antonini, S. Chatterjee, C.L. Rodriguez, M. Morscher, B. Pattabiraman, V. Kalogera, F.A. Rasio, *Astrophys. J.* 816 (2016) 65, arXiv:1509.05080.
- [5] A. Rebei, E.A. Huerta, S. Wang, S. Habib, R. Haas, D. Johnson, D. George, *Phys. Rev. D* 100 (2019) 044025, arXiv:1807.09787 [gr-qc].
- [6] E.A. Huerta, et al., *Phys. Rev. D* 100 (2019) 064003, arXiv:1901.07038 [gr-qc].
- [7] S. Habib, E.A. Huerta, arXiv e-prints, arXiv:1904.09295, 2019, arXiv:1904.09295 [gr-qc].
- [8] D. Johnson, E.A. Huerta, R. Haas, *Class. Quantum Gravity* 35 (2018) 027002, arXiv:1708.02941 [gr-qc].
- [9] E.A. Huerta, C.J. Moore, P. Kumar, D. George, A.J.K. Chua, R. Haas, E. Wessel, D. Johnson, D. Glennon, A. Rebei, A.M. Holgado, J.R. Gair, H.P. Pfeiffer, *Phys. Rev. D* 97 (2018) 024031, arXiv:1711.06276 [gr-qc].
- [10] 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, B. Szilagyi, *Phys. Rev. D* 95 (2017) 024038, arXiv:1609.05933 [gr-qc].
- [11] I. Hinder, L.E. Kidder, H.P. Pfeiffer, *Phys. Rev. D* 98 (2018) 044015, arXiv:1709.02007 [gr-qc].
- [12] Z. Cao, W.-B. Han, *Phys. Rev. D* 96 (2017) 044028, arXiv:1708.00166 [gr-qc].
- [13] T. Hinderer, S. Babak, *Phys. Rev. D* 96 (2017) 104048, arXiv:1707.08426 [gr-qc].
- [14] C.J. Hailey, et al., *Nature* 556 (2018), <https://doi.org/10.1038/nature25029>.
- [15] A.C. Sippel, J.R. Hurley, *Mon. Not. R. Astron. Soc.* 430 (2013) L30, arXiv:1211.6608 [astro-ph.GA].
- [16] J. Strader, L. Chomiuk, T.J. Maccarone, J.C.A. Miller-Jones, A.C. Seth, *Nature* 490 (2012) 71, arXiv:1210.0901 [astro-ph.HE].
- [17] J. Samsing, D.J. D'Orazio, A. Askar, M. Giersz, arXiv e-prints, arXiv:1802.08654 [astro-ph.HE], 2018.
- [18] L. Randall, Z.-Z. Xianyu, *Astrophys. J.* 853 (2018) 93, arXiv:1708.08569 [gr-qc].
- [19] J. Samsing, T. Ilan, *Mon. Not. R. Astron. Soc.* 476 (2018) 1548, arXiv:1706.04672 [astro-ph.HE].
- [20] E.A. Huerta, S.T. McWilliams, J.R. Gair, S.R. Taylor, *Phys. Rev. D* 92 (2015) 063010, arXiv:1504.00928 [gr-qc].
- [21] J. Samsing, T. Ilan, *Mon. Not. R. Astron. Soc.* 482 (2019) 30, arXiv:1709.01660 [astro-ph.HE].
- [22] E.A. Huerta, P. Kumar, S.T. McWilliams, R. O'Shaughnessy, N. Yunes, *Phys. Rev. D* 90 (2014) 084016, arXiv:1408.3406 [gr-qc].
- [23] J. Samsing, D.J. D'Orazio, *Mon. Not. R. Astron. Soc.* 481 (2018) 5445, arXiv:1804.06519 [astro-ph.HE].
- [24] D.J. D'Orazio, J. Samsing, *Mon. Not. R. Astron. Soc.* 481 (2018) 4775, arXiv:1805.06194 [astro-ph.HE].
- [25] C.L. Rodriguez, P. Amaro-Seoane, S. Chatterjee, K. Kremer, F.A. Rasio, J. Samsing, C.S. Ye, M. Zevin, *Phys. Rev. D* 98 (2018) 123005, arXiv:1811.04926 [astro-ph.HE].
- [26] K. Kremer, C.L. Rodriguez, P. Amaro-Seoane, K. Breivik, S. Chatterjee, M.L. Katz, S.L. Larson, F.A. Rasio, J. Samsing, C.S. Ye, M. Zevin, arXiv:1811.11812, 2018, arXiv:1811.11812 [astro-ph.HE].
- [27] J. Lopez, Martin, A. Batta, E. Ramirez-Ruiz, I. Martinez, J. Samsing, arXiv:1812.01118, 2018, arXiv:1812.01118 [astro-ph.HE].
- [28] B.-M. Hoang, S. Naoz, B. Kocsis, F.A. Rasio, F. Dosopoulou, *Astrophys. J.* 856 (2018) 140, arXiv:1706.09896 [astro-ph.HE].
- [29] L. Randall, Z.-Z. Xianyu, *Astrophys. J.* 864 (2018) 134, arXiv:1802.05718 [gr-qc].
- [30] F. Antonini, F.A. Rasio, *Astrophys. J.* 831 (2016) 187, arXiv:1606.04889 [astro-ph.HE].
- [31] E.A. Huerta, D.A. Brown, *Phys. Rev. D* 87 (2013) 127501, arXiv:1301.1895 [gr-qc].
- [32] J. Samsing, E. Ramirez-Ruiz, *Astrophys. J. Lett.* 840 (2017) L14, arXiv:1703.09703 [astro-ph.HE].
- [33] N.W.C. Leigh, A.M. Geller, B. McKernan, K.E.S. Ford, M.M. Mac Low, J. Bellovary, Z. Haiman, W. Lyra, J. Samsing, M. O'Dowd, B. Kocsis, S. Endlich, *Mon. Not. R. Astron. Soc.* 474 (2018) 5672, arXiv:1711.10494 [astro-ph.GA].
- [34] J. Samsing, A. Askar, M. Giersz, *Astrophys. J.* 855 (2018) 124, arXiv:1712.06186 [astro-ph.HE].
- [35] B. Farr, D.E. Holz, W.M. Farr, *Astrophys. J.* 854 (2018) L9, arXiv:1709.07896 [astro-ph.HE].
- [36] N. Fernandez, S. Profumo, *J. Cosmol. Astropart. Phys.* 1908 (2019) 022, arXiv:1905.13019 [astro-ph.HE].
- [37] B. Abbott, R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. Adhikari, V. Adya, C. Affeldt, et al., *Phys. Rev. X* 9 (2019), <https://doi.org/10.1103/physrevx.9.031040>.
- [38] B.P. Abbott, R. Abbott, T.D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R.X. Adhikari, V.B. Adya, C. Affeldt, et al., *Astrophys. J.* 892 (2020) L3.
- [39] R. Abbott, et al., LIGO Scientific, Virgo, arXiv:2004.08342 [astro-ph.HE], 2020.
- [40] D. George, E.A. Huerta, *Phys. Rev. D* 97 (2018) 044039, arXiv:1701.00008 [astro-ph.IM].
- [41] D. George, E.A. Huerta, *Phys. Lett. B* 778 (2018) 64, arXiv:1711.03121 [gr-qc].
- [42] D. George, E.A. Huerta, arXiv e-prints, arXiv:1711.07966 [gr-qc], 2017.
- [43] H. Gabbard, M. Williams, F. Hayes, C. Messenger, *Phys. Rev. Lett.* 120 (2018) 141103, arXiv:1712.06041 [astro-ph.IM].
- [44] X. Fan, J. Li, X. Li, Y. Zhong, J. Cao, *Sci. China, Phys. Mech. Astron.* 62 (2019) 969512, arXiv:1811.01380 [astro-ph.IM].
- [45] T.D. Gebhard, N. Kilbertus, I. Harry, B. Schölkopf, *Phys. Rev. D* 100 (2019) 063015.
- [46] H. Shen, D. George, E.A. Huerta, Z. Zhao, in: ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP, 2019, pp. 3237–3241.
- [47] W. Wei, E.A. Huerta, *Phys. Lett. B* 800 (2020) 135081, arXiv:1901.00869 [gr-qc].
- [48] A. Torres-Forné, E. Cuoco, J.A. Font, A. Marquina, arXiv e-prints, arXiv:2002.11668, 2020, arXiv:2002.11668 [gr-qc].
- [49] E.A. Huerta, G. Allen, I. Andreoni, J.M. Antelis, E. Bachelet, G.B. Berriman, F.B. Bianco, R. Biswas, M. Carrasco Kind, K. Chard, M. Cho, P.S. Cowperthwaite, Z.B. Etienne, M. Fishbach, F. Forster, D. George, T. Gibbs, M. Graham, W. Gropp, R. Gruendl, A. Gupta, R. Haas, S. Habib, E. Jennings, M.W.G. Johnson, E. Katsavounidis, D.S. Katz, A. Khan, V. Kindratenko, W.T.C. Kramer, X. Liu, A. Mahabal, Z. Marka, K. McHenry, J.M. Miller, C. Moreno, M.S. Neubauer, S. Oberlin, A.R. Olivas, D. Petravick, A. Rebei, S. Rosofsky, M. Ruiz, A. Saxton, B.F. Schutz, A. Schwing, E. Seidel, S.L. Shapiro, H. Shen, Y. Shen, L.P. Singer, B.M. Sipocz, L. Sun, J. Towns, A. Tsokaros, W. Wei, J. Wells, T.J. Williams, J. Xiong, Z. Zhao, *Nat. Rev. Phys.* 1 (2019) 600, arXiv:1911.11779 [gr-qc].
- [50] T.S. Yamamoto, T. Tanaka, Use of conditional variational auto encoder to analyze ringdown gravitational waves, arXiv:2002.12095 [gr-qc], 2020.
- [51] H. Gabbard, C. Messenger, I.S. Heng, F. Tonolini, R. Murray-Smith, Bayesian parameter estimation using conditional variational autoencoders for gravitational-wave astronomy, arXiv:1909.06296 [astro-ph.IM], 2019.
- [52] S.R. Green, C. Simpson, J. Gair, Gravitational-wave parameter estimation with autoregressive neural network flows, arXiv:2002.07656 [astro-ph.IM], 2020.
- [53] V. Varma, S.E. Field, M.A. Scheel, J. Blackman, L.E. Kidder, H.P. Pfeiffer, *Phys. Rev. D* 99 (2019) 064045.
- [54] L.D. Dalcín, R.R. Paz, P.A. Kler, A. Cosimo, New computational methods and software tools, *Adv. Water Resour.* 34 (2011) 1124.
- [55] L. Dalcín, R. Paz, M. Storti, J. D'Elia, *J. Parallel Distrib. Comput.* 68 (2008) 655.
- [56] L. Dalcín, R. Paz, M. Storti, *J. Parallel Distrib. Comput.* 65 (2005) 1108.
- [57] National Center for Supercomputing Applications, Illinois Campus Cluster, 2020, <https://campuscluster.illinois.edu/>.
- [58] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, K. Kavukcuoglu, Wavenet: a generative model for raw audio, arXiv:1609.03499 [cs.SD], 2016.
- [59] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, a. Graves, N. Kalchbrenner, A. Senior, K. Kavukcuoglu, in arXiv (2016).
- [60] T. Damour, *Phys. Rev. D* 64 (2001) 124013, arXiv:gr-qc/0103018 [gr-qc].
- [61] B.P. Abbott, et al., LIGO Scientific, Virgo, *Phys. Rev. Lett.* 116 (2016) 241102, arXiv:1602.03840 [gr-qc].
- [62] M. Pürrer, M. Hannam, F. Ohme, *Phys. Rev. D* 93 (2016) 084042, arXiv:1512.04955 [gr-qc].
- [63] National Center for Supercomputing Applications, Hardware-Accelerated Learning (HAL) Deep Learning Cluster, 2018, <https://wiki.ncsa.illinois.edu/display/ISL20/HAL+cluster>.
- [64] E.A. Huerta, A. Khan, E. Davis, C. Bushell, W.D. Gropp, D.S. Katz, V. Kindratenko, S. Koric, W.T.C. Kramer, B. McGinty, K. McHenry, A. Saxton, arXiv:2003.08394, 2020, arXiv:2003.08394 [physics.comp-ph].
- [65] N.S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, P.T.P. Tang, On large-batch training for deep learning: generalization gap and sharp minima, arXiv:1609.04836 [cs.LG], 2016.
- [66] Y. You, J. Li, S. Reddi, J. Hseu, S. Kumar, S. Bhojanapalli, X. Song, J. Demmel, K. Keutzer, C.-J. Hsieh, Large batch optimization for deep learning: training BERT in 76 minutes, arXiv:1904.00962 [cs.LG], 2019.
- [67] A. Khan, E.A. Huerta, A. Das, Deep Learning Hub, 2020.
- [68] B. Blaiszik, L. Ward, M. Schwarting, J. Gaff, R. Chard, D. Pike, K. Chard, I. Foster, *MRS Commun.* 9 (2019) 1125–1133.
- [69] R. Chard, Z. Li, K. Chard, L. Ward, Y. Babuji, A. Woodard, S. Tuecke, B. Blaiszik, M.J. Franklin, I. Foster, in: 2019 IEEE International Parallel and Distributed Processing Symposium, IPDPS, 2019, pp. 283–292.
- [70] Asad Khan, Physics-Inspired Deep Learning for Gravitational Wave Astrophysics, 2020, https://khanx169.github.io/smr_bbm_v2/interactive_results.html.