

Source-Agnostic Gravitational-Wave Detection with Recurrent Autoencoders

Eric A. Moreno , Jean-Roch Vlimant , Maria Spiropulu 

California Institute of Technology, Pasadena, California 91125, USA

Bartłomiej Borzyszkowski  ‡

Gdansk University of Technology, Narutowicza 11/12, 80-233, Gdansk, Poland

Maurizio Pierini 

European Organization for Nuclear Research, 1211 Geneva 23, Switzerland

E-mail: maurizio.pierini@cern.ch

July 2021

Abstract. We present an application of anomaly detection techniques based on deep recurrent autoencoders to the problem of detecting gravitational wave signals in laser interferometers. Trained on noise data, this class of algorithms could detect signals using an unsupervised strategy, i.e., without targeting a specific kind of source. We develop a custom architecture to analyze the data from two interferometers. We compare the obtained performance to that obtained with other autoencoder architectures and with a convolutional classifier. The unsupervised nature of the proposed strategy comes with a cost in terms of accuracy, when compared to more traditional supervised techniques. On the other hand, there is a qualitative gain in generalizing the experimental sensitivity beyond the ensemble of pre-computed signal templates. The recurrent autoencoder outperforms other autoencoders based on different architectures. The class of recurrent autoencoders presented in this paper could complement the search strategy employed for gravitational wave detection and extend the reach of the ongoing detection campaigns.

Keywords: Machine Learning, Unsupervised Learning, Anomaly Detection

‡ Also at CERN and Intel Technology Gdansk, Poland.

1. Introduction

The detection of gravitational waves (GW) from stellar binaries such as black hole and neutron star mergers have ushered in a new era of analyzing the Universe. Operating together, the Laser Interferometer Gravitational-wave Observatory (LIGO) [1] and the Virgo Interferometer [2] can peer into deep space giving astronomers the ability to uncover and localize stellar processes through their gravitational signature. The first observation of a binary black hole merger (GW150914) [3] has given way to a plethora of GW events, and notably to the observation of intermediate-size black holes [4] and neutron star mergers [5], an event that marked the beginning of the multi-messenger astronomy era [6].

Instrumental on the software side of these observations are the algorithms which identify the faint GW signals in an environment characterized by overwhelming classical and quantum noise. The most sensitive detection algorithm, Matched Filtering (MF) [7], consists of matching incoming data with templates of simulated GW shapes, covering the parameter space of binary masses [8, 9, 10, 11, 12], which is then used to identify the signal. At the same time, they offer an estimate of the astrophysical parameters associated to the detected GW event, such as the nature and mass of the two merging objects. This method was extremely functional to the success of the LIGO and VIRGO observation campaigns. On the other hand, by relying on pre-computed templates, it could lead to missing events for which a template is not available. This could be an event originating from computationally prohibitive conditions [13, 14, 15, 16], or even some sort of unforeseen GW source.

Detecting GWs is certainly one of the hardest challenges faced in fundamental science in the recent years. Given how weak a GW signal is when compared to typical noise levels, it is natural to look at Machine Learning (ML), and especially to Deep Learning (DL), in order to improve a signal detection capability. For instance, Refs. [17, 18] discusses how Convolutional 1D networks [19, 20] can be trained to extract a variety of GW signals from highly noisy data. This is an example of a supervised classifier, i.e., a classifier trained to separate different populations in data (e.g., signal vs. noise) by matching a given set of ground-truth labels. While classifiers could certainly contribute to enhance the current state-of-the-art detection capability, they rely on pre-defined signal much like the MF technique. In other words, they are designed to possibly improve the detection accuracy but they are not necessarily going to extend the detector sensitivity to exotic signals outside the portfolio of pre-simulated templates. In fact, these networks are typically trained on labelled data from simulation, so even in this case the capability of simulating a given signal is an underlying requirement. On the other hand, GW detection comes with the need of going beyond signal signatures that can be emulated. Besides the search for exotic sources, model independent strategies could be useful to deal with practical issues such as glitch detection [21].

In this paper, we investigate the possibility of rephrasing the problem of GW detection as an anomaly detection task. By anomaly detection we mean the use of

a one-class classifier, in this case a Deep Autoencoder (AE) [22], to identify outlier populations in an unlabelled dataset. An autoencoder is a compression-decompression algorithm that is trained to map a given set of inputs into itself, by first compressing the input into a point in a learned latent space (encoding) and then reconstructing it from the encoded information (decoding).

Once trained on standard events, the autoencoder might fail to reconstruct samples of different kind (the anomalies). Any input-to-output distance measurement can then be used to identify these anomalies. Under the assumption that these anomalies are rare, one can directly train these AEs on data, looking for the set of AE parameter values that minimize the difference between the input and the output, using some distance D as a loss function. By taking as a reference the distribution of D on data, one can label anomalies as the outlier events of this distribution. This very same approach was recently discussed in Ref. [23], where the reach of Convolutional autoencoders for GW detection is investigated. In this work, we consider two recurrent AE architectures: Long-short memory networks (LSTMs) [24], and Gated Recurrent Units (GRUs) [25]). For comparison, we consider alternative AE architectures: dense (i.e., fully connected) neural networks (DNNs) and Convolutional Neural Networks (CNNs).

The main advantage of this unsupervised strategy is that one algorithm is potentially sensitive to multiple signal typologies. On the other hand, this gain in flexibility is typically followed by a loss in accuracy. For a specific signal, an algorithm trained with an unsupervised procedure on unlabeled data is typically less accurate than a supervised classifier trained on labeled data.

The proposed strategy comes with another remarkable advantage: under the assumption that anomalies are rare instances in the processed data, autoencoders can be trained directly on data, without relying on signal or noise simulation. Instead, supervised algorithms require labels which, in the case of rare signals like those considered in this study, are obtained using synthetic data (e.g., from Monte Carlo simulation). Assuming this training could happen in real time, the AE could adapt to changing experimental conditions and limit the occurrence of false claims.

This paper is organized as follows: related works describing ML approaches to GW detection are briefly discussed in Section 2. The dataset utilized for this study is described in Section 3. The autoencoder architectures are described in Section 4, together with the corresponding classifiers used to benchmark performance. Results are presented in Section 5. Conclusions are given in Section 6.

2. Related work

Besides MFs, Bayesian inference libraries [26, 27] have been created to estimate the properties of a gravitational-wave source. Unsupervised anomaly detection of transients has also been proposed using a temporal version of k-nearest neighbors (kNN) [28].

Most DL approaches for GW classification involve supervised learning techniques, which typically provide competitive accuracy by exploiting network non-linearity and

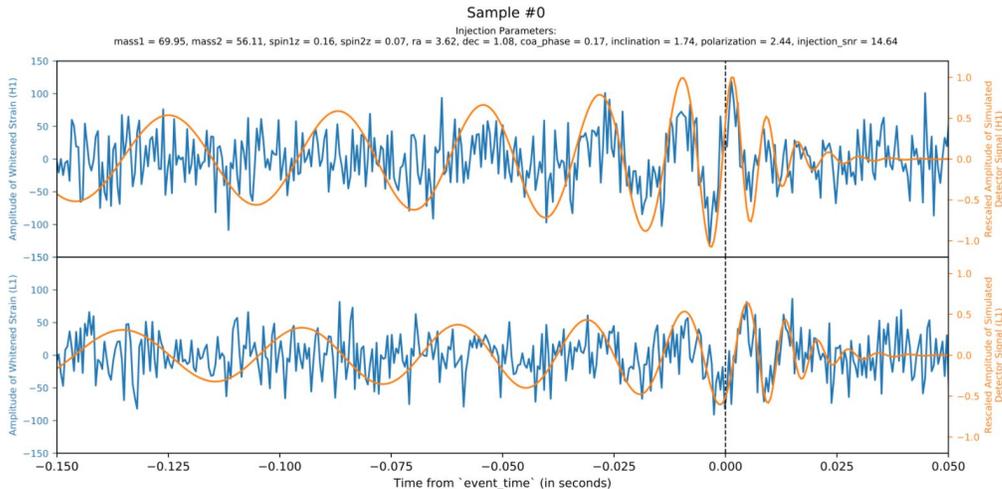


Figure 1. Resulting strains from simulated gravitational waves on the Livingston (L1) and Hanford (H1) detectors from the coalescence of a $70 M_{\odot}$ and a $56 M_{\odot}$ black holes, obtained using the GGWD package [34]. The re-scaled amplitude of the GW signal is illustrated in orange and the strain on the detector including this signal is illustrated in blue.

the information provided by ground-truth labels.

Principal Component Analysis [29] (PCA), which performs a linear orthogonal transformation of a set of (possibly correlated) variables into a set of linearly uncorrelated variables, has also been introduced for transient detection [30, 31]. This can give a quick characterization of the intrinsic properties of a data sample. The DL methods discussed in this paper are generalizations of this approach to include nonlinear compression.

Boosting Neural Networks [32], which use a combination of unsupervised and supervised learning techniques, have also been used for GW classification [33]. This method performs an unsupervised hierarchical clustering on the incoming data to identify possible groups and a supervised Bayesian classifier to do the final classification. BNNs can classify a number of different injections including Gaussian, Ringdowns, Supernovae, white noise bursts, mergers, etc. Importantly, the architecture requires sufficient statistics to cluster in an unsupervised manner, which makes it not ideal for the rare processes that we focus on in this paper.

Finally, CNNs have achieved high accuracy while having the added benefit of parameter estimation to infer useful parameters of the GWs such as masses and spins of the binary merger components [17, 18]. In a similar approach to the one discussed in this paper, Ref. [23] discusses the use of CNN autoencoders in order to classify GWs. For comparison, the same architecture as specified in [23] is implemented in Section 4 alongside the recurrent autoencoders.

3. Data samples

This study is performed on a sample of synthetic data, generated for this study and available on Zenodo [35, 36]. Data are generated using the GGWD library [34].

Noise events occur when no signal is overlapped to the detector noise. Detector noise is generated at a specified power spectral density (PSD), which is computed as a function of the length of noise, time step of the noise, and noise weighting to color the noise. This is done using PyCBC [37]. This approach to simulated data generation ignores glitches, blips, and other transient sources of detector noise §.

In absence of exotic signal sources, we use traditional GW signals to assess the detection performance. We consider two kinds of GW sources: Binary black hole (BBH) and Binary neutron star (BNS) mergers. Signal events are generated simulating GW production from compact binary coalescences using PyCBC [37], which itself uses algorithms from LIGO’s LAL Suite [38]. Signal event containing GWs were created overlaying simulated GWs on top of detector noise. This provides an analogous situation to a real GW, in which the strain from the incoming wave is recorded in combination with the normal detector noise.

The dataset consists of 300,000 noise samples. Each sample corresponds to 8 seconds of data, sampled at 2048 Hz. We consider a LIGO-like experimental setup, with two detectors (L1 and H1) taking data simultaneously at different locations. For each detector, a sample is represented as a one-dimensional array with 16,384 entries.

Figure 1 shows one of the simulated signal events: a coalescence of two black holes, with masses set to 70 and 56 stellar masses (M_{\odot}), detected by the L1 and H1 detectors. The generation parameters are listed in the figure: the spin of the two black holes, right ascension, declination, coalescence phase, inclination, polarization, and signal-to-noise ratio [34]. In absence of any source of noise, the signal would appear as shown by the orange line. Once the noise is added, the detectable signal corresponds to the blue line.

This dataset is split in three parts: 192,000 training samples (64%), 48,000 validation samples (16%), and 60,000 test samples (20%). The training and validation datasets are used in the optimal-parameter learning process, while the test dataset is used to assess the algorithm performance after training, together with the signal samples.

The BBH sample is generated as follows:

- SEOBNRv4 [39] Approximant.
- Masses independently and uniformly varied within [10, 80] M_{\odot} .
- Spins independently and uniformly varied within [0, 0.998].
- Injection signal-to-noise (SNR) ratio uniformly varied within [5, 20].

§ For this reason, the considered AEs could be re-purposed for anomaly detection algorithms to identify detector glitches. The main difference between these kind of glitch anomalies and those of astronomical significance would be the lack of coincidence of anomalies across different detectors, a clear indication of a stellar anomalous signal.

- Coalescence phase uniformly varied within $[0, 2\pi]$.
- Inclination uniformly varied within $[0, \pi]$.

The BNS sample is generated as follows:

- IMRPhenomDNRTidal_v2 [40] Approximant.
- Masses independently and uniformly varied within $[5, 10] M_{\odot}$.
- Spins independently and uniformly varied within $[0, 0.998]$.
- Injection signal-to-noise (SNR) ratio uniformly varied within $[5, 20]$.
- Coalescence phase uniformly varied within $[0, 2\pi]$.
- Inclination uniformly varied within $[0, \pi]$.

Data are whitened with a Fast Fourier Transform integration length of 4 seconds and a duration of the time-domain Finite Impulse Filter whitening filter of 4 seconds to remove the underlying correlation in the data [41]. Then, band-pass filtering was applied to remove high frequency (above 250 Hz) and low frequency (below 50 Hz) components from the data. Doing so, background from outside the current interferometer sensitivity range is discarded. To facilitate the data processing and learning by the network, the data are scaled absolutely to a $[0, 1]$ range. Each 8 second event is then cropped to 2.5 sec around the GW event, with the GW happening at a random time after the beginning of the time window. The delay is uniformly sampled in a $[0.5, 2]$ sec interval. This choice allows us to take into account the appropriate time for the signal ring-up/ring-down. This is done to assure that the model classifiers are not biased to a certain time period within the event, which would occur if the simulated data has GWs appearing at only a single time within the event window.

4. Network architectures

Autoencoders are algorithms that project an input sample $x \in \mathcal{X}$ to its encoded projection z in a latent space \mathcal{Z} , typically of lower dimension than the input space. The encoded projection z is then decoded to a reconstructed $\hat{x} \in \mathcal{X}$. The network parameters determine how x is projected to z and then back to \hat{x} . Their values are fixed minimizing some input-to-output distance, used as a loss function in the network training. In this study, we consider the mean-squared error (MSE) between each element of the input array and the corresponding output. We consider several network architectures, all structured according to a common scheme with the decoder mirroring as close as possible the encoder architecture. Three specific architectures are introduced, with CNN, LSTM, or GRU layers.

As a comparison to recurrent layers, we implement a CNN AE similar to Ref. [23] with an input of five one-dimensional images of shape $(1024, 1)$ corresponding to a 2.5-second interval sampled at 2048 Hz. The encoder consists of two one-dimensional convolutional layers with filters of size $[256, 128]$ and kernel of size 3, coupled with a maxpool layer of size 2. The decoder mirrors this architecture with an upsampling layer

of size 2, and two one-dimensional convolutional layers with filter size [256, 1] and kernel of size 3. In addition, DNN AEs with number of nodes 100, 50, 10 and 10, 50, 100 were attempted but yielded worse results. This is expected as DNN AEs are not specialized in time-series data, unlike recurrent architectures.

The LSTM and GRU networks function similarly, utilizing LSTM and GRU cells instead of simple CNN/DNN nodes. The latent space bottleneck in this representation is created by instructing the LSTM (or GRU) cells to only return their final state in the encoding phase and then repeating that final state as a vector which can then be input to the decoder. The input to these recurrent architectures are 51 one-dimensional arrays of shape (100, 1) corresponding to a 2.5-second interval sampled at 2048 Hz. The encoder consists of 2 layers, with number of units 32 and 8, which produces a latent representation by only outputting the final LSTM state vector on the bottleneck layer. The decoder consists of two layers with 8, 32 nodes, which is then multiplied by a final temporal slice of a dense layer, yielding the same dimensions as the input representation. For illustration, the architecture of the LSTM model is shown in Fig. 2.

The CNN, LSTM, and GRU autoencoders are all trained on unlabeled detector noise data, with no introduction to signal distributions. Doing so, the latent space representation returned by the encoder is exclusively a function of detector noise. As a result, the MSE error is relatively consistent during periods with exclusively noise, but it might instantaneously increase when a signal event passes through the autoencoder. An example of such a spike in the MSE loss as a function of time is shown in Fig. 3. The spike is typically due to the fact that the encoding/decoding sequence learned on noise might not be optimal for a previously unobserved kind of input data. As a result, the distance between the input and the output could be larger for anomalous data, up to generate a spike. Operationally, one could then monitor the MSE value returned by the algorithm, and the detection of a signal could be correlated to the observation of a spike above threshold.

The performance of the three AE models is assessed comparing their accuracy on benchmark signal samples to that obtained from binary CNN classifiers, trained on the same data and the corresponding labels. Different classifiers are trained for different signal topologies. The classifier architecture is loosely equivalent to that of the CNN encoder from Ref. [17, 18]. In particular, it consists of four convolution layers, with 64, 128, 256, and 512 filters respectively, and two fully connected layers with 128 and 64 nodes, respectively. A ReLU activation function was used throughout. Kernel sizes of 16x16, 16x16, 16x16, and 32x32 with a stride of 1 for the convolutional layers and 4x4 with a stride of 4 for all the max pooling layers were used with dilations of 1, 2, 2, and 2 in the corresponding convolutional layers. A sigmoid function is used for the single-node output layer. An LSTM-implementation of the supervised classifier was also attempted but yielded results far worse than the CNN classifier method, so it is not included in this study.

Two classifiers of this kind are trained, using a dataset of 300,000 samples, consisting an equal fraction of noise events and one of the two classes of signal (BBH and BNS)

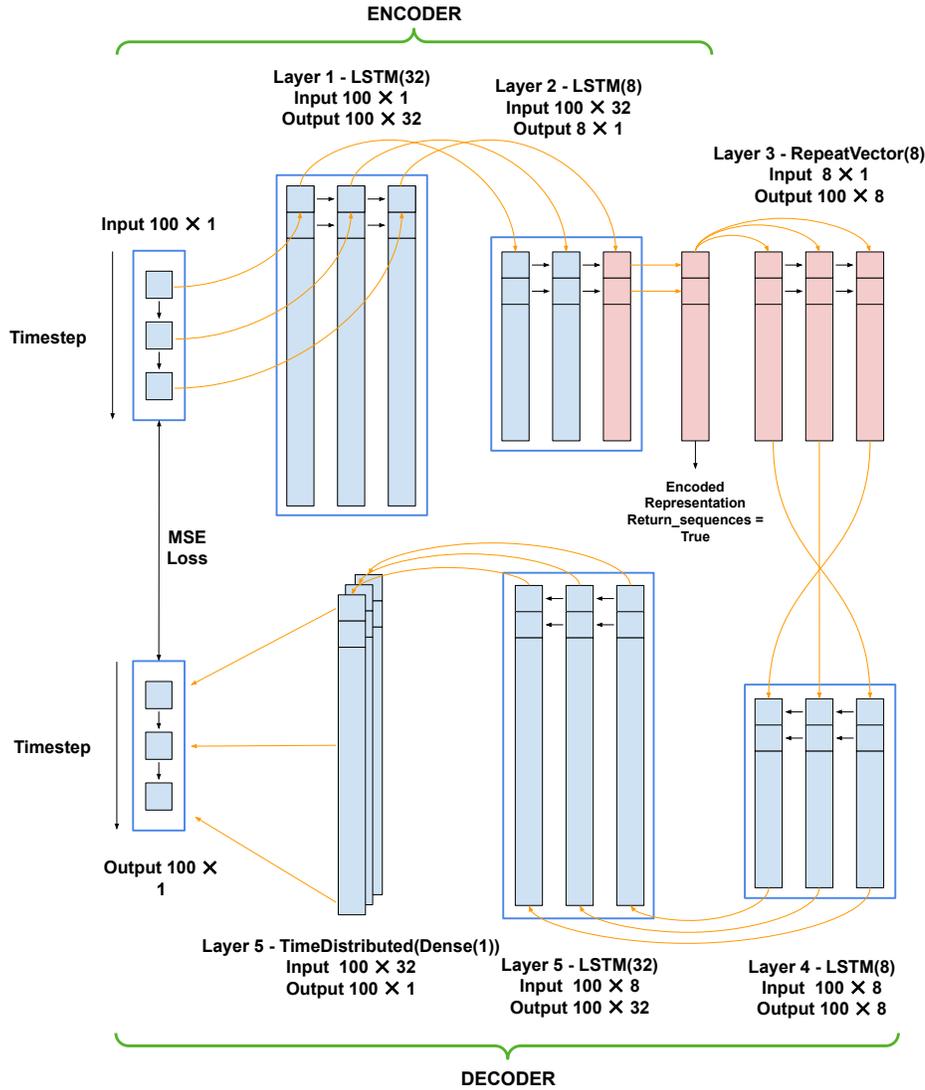


Figure 2. Graphical illustration of LSTM autoencoder. The general encoder structure still remains, and the latent space is created by setting `return_sequences=False` in the LSTM cell, passing through only the final time step of the LSTM. This latent space is then repeated using a `RepeatVector`, forcing the encoder to create some important latent representation as its final output which can then be sent to the decoder.

considered in this study. The training is performed minimizing a binary cross entropy error loss function on the training sample of Sec. 3, using the validation set to optimize the training and the test set to evaluate the model performance.

The classifier is tested on noise samples, as well as on BBH and BNS events. When tested on the same kind of signal it is trained on, the classifier accuracy is used to estimate the best accuracy that the AE could reach and, consequently, the loss in accuracy due to the use of an unsupervised approach. When testing the classifier on the signal it was not trained on, we can instead compare the generalization property of the autoencoder to that

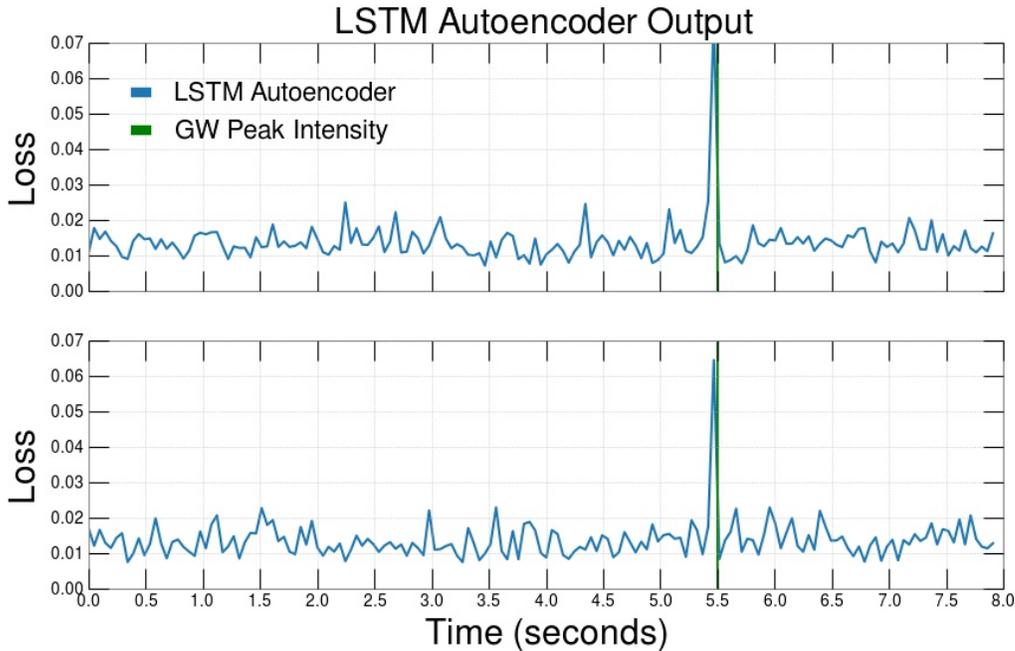


Figure 3. Spike/dip in MSE loss for an LSTM autoencoder over two 8-second event sampled at 2048 Hz (split into 100-timestep inputs) signaling the detection of a gravitational wave. The AE models are never presented with any GW signals (green) during training. Detection thresholds are set by fixing a FPR, which is done continuously to create a ROC curve in Section 5.

of a supervised algorithm. The two tests provide an assessment of the balance between accuracy and generalization power and demonstrate the complementarity between our approach and a standard template-based method. In practice, one could implement as many supervised algorithms as known GW sources, while using an unsupervised algorithm to be sensitive to unexpected signal sources (and non-coincident signals across multiple interferometers, in cases of glitch detection and data quality monitoring).

5. Results

Figure 4 shows the receiver operating characteristic (ROC) curves for three autoencoder architectures (LSTM, GRU, CNN). The curves are obtained considering a single detector, i.e., no coincidence is enforced at this stage. In the left (right) figure, the ROC curves are evaluated on noise and a signal sample of BBH (BNS) merger data. For comparison, the CNN classifiers trained on both datasets are shown.

As the ROC curves show, the LSTM architecture provides the best accuracy among the AEs. The LSTM AE accuracy is worse than that of the classifier trained on the correct signal hypothesis, but better than that of the classifier trained on the opposing signal hypothesis. The performance comparison is quantified in Table 1, where the false positive rates (FPRs) corresponding to fixed values of the true positive rates (TPRs)

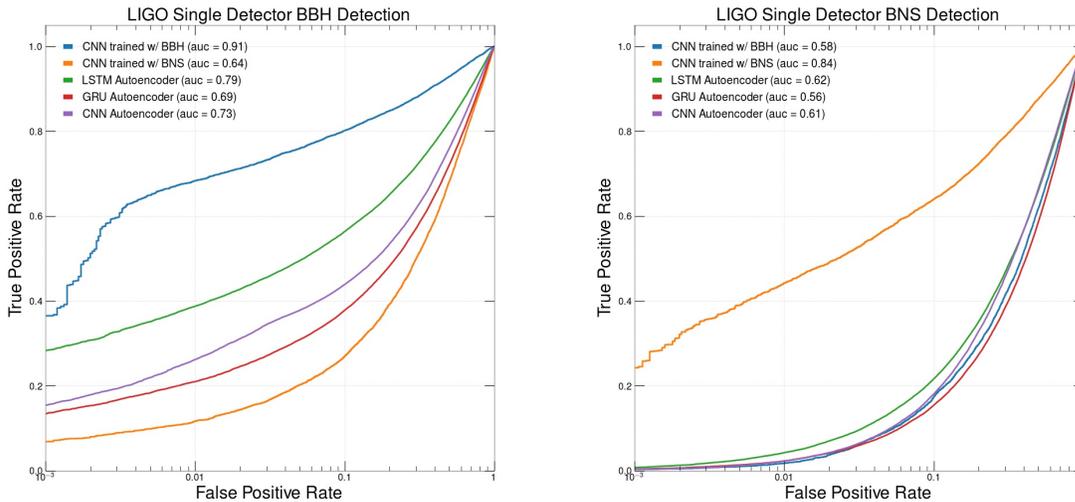


Figure 4. Single-detector ROC curves for the LSTM, GRU, CNN Autoencoders compared to the corresponding CNN supervised architecture. The figure shows the ROC curves for supervised CNNs trained on the BBH and BNS merger data, when the inference is performed on noise and a signal sample of BBH (left) or BNS (right) events.

are shown. The ROC curves and the values shown in Table 1 quantify the trade-off between accuracy and generalization that motivates this study. This makes autoencoders especially useful to potentially discover unexpected GW sources, as well as GW sources that cannot be modeled by traditional simulation techniques.

Table 1. True-positive rates for BBH and BNS merger detection at 10% and 1% false-positive rates, for autoencoders trained on noise and for binary CNN classifiers trained on BBH and BNS simulations. The autoencoder architecture with the best unsupervised results is marked in bold.

BBH signal vs. noise					
FPR	LSTM-AE	GRU-AE	CNN-AE	BBH-CNN	BNS-CNN
0.1	56.4%	37.8%	43.8%	80.1%	26.9%
0.01	38.7%	21.0%	26.2%	68.3%	11.7%
BNS signal vs. noise					
FPR	LSTM-AE	GRU-AE	CNN-AE	BBH-CNN	BNS-CNN
0.1	21.7%	15.4%	16.6%	17.3%	64.0%
0.01	4.2%	2.3%	2.3%	1.75%	44.2%

This advantage is especially marked with BBHs, which have larger masses and thus larger SNRs. Thus, autoencoder architectures will likely have an advantage with higher mass-range mergers in the regime where supervised learning models cannot generalize but which still have large SNRs. This is as opposed to BNSs, which have lower mass values and consequently lower SNR signatures. In this case, the generalization performance

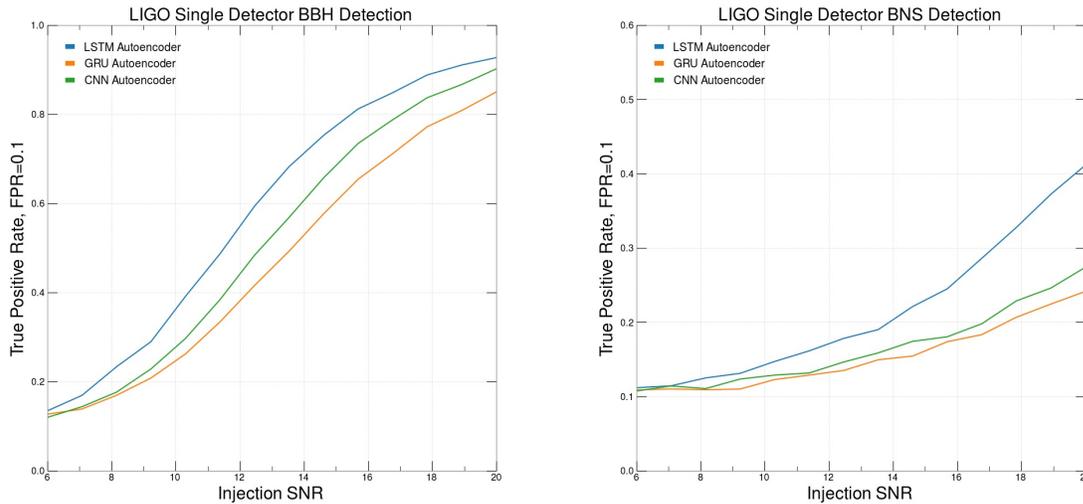


Figure 5. True positive rates for the LSTM Autoencoder at a fixed False Positive Rate (FPR=0.1) with BBH (left) and BNS (right) events at variable Signal-to-noise ratios.

stagnates for both of the models, meaning that both models are extracting the same amount of signal out of the events. Still, the autoencoder models perform better than the supervised algorithm trained on the wrong signal hypothesis.

The TPR values quoted on Table 1 are obtained averaging across the SNR, which is uniformly distributed in a $[5, 20]$ range. On the other hand, the TPRs of the AE models depend strongly on the SNR value, as shown in Fig. 5 both for BBH and BNS merger events. As shown in the figure, the LSTM AE guarantees better performance across the considered range of SNR values. While the improvement (e.g., with respect to the CNN AE) is roughly constant for BBH events, in the case of BNS events the LSTM AE is particularly better than the other architectures for large SNR values. Overall, this study reinforces the idea that the LSTM AE is the most robust choice among those we considered.

In a realistic exploitation of this algorithm, one would define a threshold above which the data would be called a potential signal. Doing so, one would like to keep the FPR at a manageable rate, while retaining a reasonable TPR value. For instance, an FPR of 10^{-4} would correspond to about one false alarm a day, low enough for a post detection assessment of the nature of the anomaly. Similarly, a FPR of 10^{-6} would correspond to about one false alarm every three months, low enough for the algorithm to be used in a real-time data processing, e.g., to serve as a trigger for multi-messenger astronomy.

The key ingredient to reach low FPR values is the exploitation of signal coincidence across multiple detectors. Since the noise across detectors is uncorrelated, single-detector FPR of $\sim 10^{-2}$ (as in Table 1) would give a global FPR of $\sim 10^{-4}$ when two instruments are put in coincidence. Clearly, the presence of uncorrelated noise overlapped to the

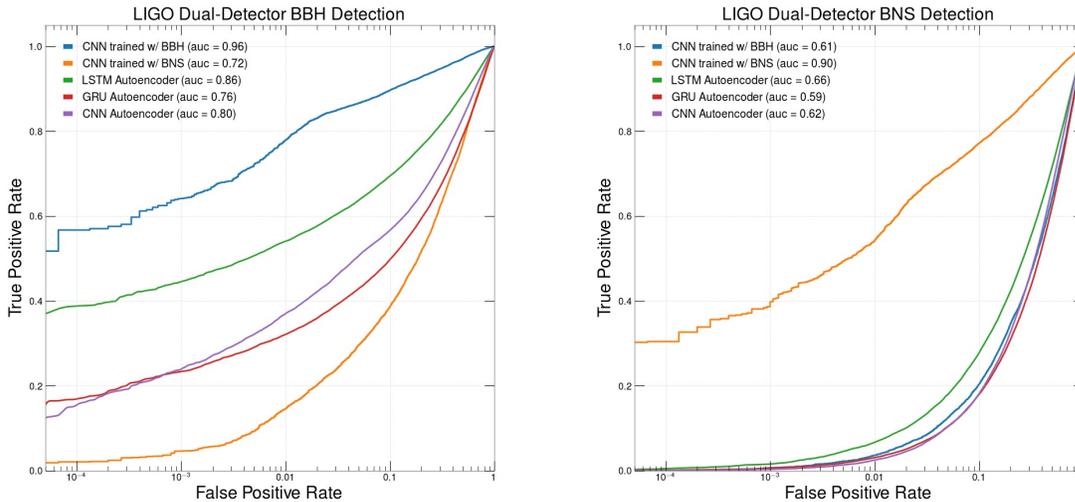


Figure 6. ROC curves for the LSTM, GRU, CNN Autoencoders, obtained by exploiting signal coincidence on two detectors. The loss is computed independently on each detector and the same threshold is applied. The test is performed on noise and a signal sample of BBH (left) or BNS (right) events.

signal dilutes the correlation of the anomaly across different devices. On the other hand, a certain level of correlation is retained. For instance, we observe a 30% correlation on the LSTM anomaly score for BBH merger events. For comparison, a 60% correlation is observed for the CNN classifier. Coincidence can be enforced requiring that two signals above a certain threshold are detected at the same time. Alternatively, one could apply a threshold on the sum of the two losses, with the idea that an MSE loss function is loosely related to the negative log likelihood, so that the sum of the loss would correspond to the negative log of the likelihood products. The former approach has the advantage of requiring the two detectors to communicate only after the anomaly event in a detector is identified. This means that the data throughput to be transmitted can be kept low. On the other hand, the latter approach provides better performance and it is considered here. In this case, one would have to find solutions to mitigate the data throughput and facilitate the communication of the detectors in real time. For instance, one could run the encoder at each experiment site and transmit the compressed data, with the decoding and coincidence check happening off-site. One should keep in mind that the LSTM model can run on a Field Programmable Gate Array within $\mathcal{O}(100)$ nsec, as demonstrated in Ref. [42].

To show this, we consider the case of two detectors and we build a ROC curve requiring a signal above a threshold on the sum of the autoencoder losses. The result is shown in Fig. 6, both for BBH and BNS mergers and quantified in Table 2. Keeping as a target a FPR of 10^{-4} , one can retain a BBH TPR comparable to that of the 10^{-2} TPR of the single-detector threshold (see Table 1), while reducing the FPR by two orders of magnitude. The gain is less striking in the case of BNS: the two-detector combination

comes with a reduction in terms of TPR, but the value achieved is still better than the square of the single-detector TPRs (e.g., 0.5% for a 10^{-4} FPR, to be compared with 0.18% from squaring the value quoted in Table 1). Overall, the improvement obtained using two detectors depends on the SNR value but certainly there is an advantage in exploiting the coincidence of the signal across detectors.

Table 2. True-positive rates for BBH and BNS merger detection at single-detector 10^{-1} and 10^{-2} false-positive rates corresponding to dual-detector FPRs of 10^{-2} and 10^{-4} , obtained by exploiting signal coincidence in two detectors. The autoencoder architecture with the best unsupervised results is marked in bold.

BBH signal vs. noise					
FPR	LSTM-AE	GRU-AE	CNN-AE	BBH-CNN	BNS-CNN
0.01	54.1%	36.1%	38.3%	78.7%	14.6%
0.0001	38.9%	18.1%	17.4%	58.2%	2.9%
BNS signal vs. noise					
FPR	LSTM-AE	GRU-AE	CNN-AE	BBH-CNN	BNS-CNN
0.01	7.4%	3.1%	2.8%	3.3%	53.7%
0.0001	0.5%	0.2%	0.2%	0.0%	31.6%

6. Conclusions

We presented an unsupervised strategy to detect GW signals from unspecified sources exploiting an AE trained on noise. The AE is trained to compress input data to a low-dimension latent space and reconstruct a representation of the input from the point in the latent space. The algorithm is optimized using as a loss function a differentiable metric, quantifying the distance between input and output data. Given a trained AE, one could identify anomalous data isolating the outlier data populating the tail of the loss distribution.

We applied this strategy to a sample of synthetic data from two GW interferometers. We explore different choices for the network architecture and compare the single-detector detection capability to that of a CNN binary classifier, trained on specific signal hypotheses. We show how a recurrent AE provides the best anomaly detection performance on benchmark BBH and BNS merger events. We show the trade-off between accuracy and generalization, when using this unsupervised strategy rather than a supervised approach (here quantified through a CNN binary classifier trained on the same data and the corresponding labels).

We show how the coincidence of two detectors, both selecting anomalies at an expected FPR $\sim 10^{-2}$, would retain a TPR of 53.1% (28.7%) for BBH (BNS) signals while giving one false alarm a day which can be easily be discarded after a post-detection analysis, e.g., with more traditional GW detection strategies.

With the same FPR, one could bring the false alarm rate to about once every three month, exploiting the coincidence of three detectors. In this case, the algorithms proposed

here could be operated in the real-time as part of a trigger system for multi-messenger astronomy, in the spirit of what is discussed in Ref. [43] for real-time data analysis at the Large Hadron Collider. Considering the relatively low computational cost of such an algorithm [42] and the high impact of a potential signal detection by this algorithm, its implementation for LIGO and VIRGO would be certainly be beneficial, despite the fact that the expected detection probability cannot be guaranteed to be high for any signal source.

Acknowledgments

We are grateful to the insight and expertise of Rana Adhikari and Hang Yu from the LIGO collaboration and Elena Cuoco from the VIRGO collaboration, who guided us on a field of research which is not our own.

Part of this work was conducted at "*iBanks*", the AI GPU cluster at Caltech. We acknowledge NVIDIA, SuperMicro and the Kavli Foundation for their support of "*iBanks*".

This work was carried on as part of the 2020 CERN OpenLab Summer Student program, which was carried on in remote mode due to the COVID pandemic.

M. P. is supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant Agreement No. 772369).

E. M. is supported by the Institute for Research and Innovation in Software for High Energy Physics (IRIS-HEP) through a fellowship in Innovative Algorithms.

This work is partially supported by the U.S. DOE, Office of Science, Office of High Energy Physics under Award No. DE-SC0011925, DE-SC0019227 and DE-AC02-07CH11359.

References

- [1] J. Aasi et al. Advanced LIGO. *Class. Quant. Grav.*, 32:074001, 2015.
- [2] F. Acernese et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Class. Quant. Grav.*, 32(2):024001, 2015.
- [3] B. P. Abbott et al. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116:061102, Feb 2016.
- [4] B.P. Abbott et al. GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2. *Phys. Rev. Lett.*, 118(22):221101, 2017. [Erratum: *Phys.Rev.Lett.* 121, 129901 (2018)].
- [5] B.P. Abbott et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.*, 119(16):161101, 2017.
- [6] B.P. Abbott et al. Multi-messenger Observations of a Binary Neutron Star Merger. *Astrophys. J. Lett.*, 848(2):L12, 2017.
- [7] B. Allen et al. FINDCHIRP: An Algorithm for detection of gravitational waves from inspiraling compact binaries. *Phys. Rev. D*, 85:122006, 2012.
- [8] B.S. Sathyaprakash and S.V. Dhurandhar. Choice of filters for the detection of gravitational waves from coalescing binaries. *Phys. Rev. D*, 44:3819–3834, Dec 1991.

- [9] R. Balasubramanian, B.S. Sathyaprakash, and S.V. Dhurandhar. Gravitational waves from coalescing binaries: Detection strategies and monte carlo estimation of parameters. *Phys. Rev. D*, 53:3033–3055, Mar 1996.
- [10] B.J. Owen. Search templates for gravitational waves from inspiraling binaries: Choice of template spacing. *Phys. Rev. D*, 53:6749–6761, Jun 1996.
- [11] T. Cokelaer. Gravitational waves from inspiraling compact binaries: Hexagonal template placement and its efficiency in detecting physical signals. *Phys. Rev. D*, 76:102004, Nov 2007.
- [12] R. Smith et al. Fast and accurate inference on gravitational waves from precessing compact binaries. *Phys. Rev. D*, 94(4):044031, 2016.
- [13] S. Klimentenko et al. Method for detection and reconstruction of gravitational wave transients with networks of advanced detectors. *Phys. Rev. D*, 93:042004, Feb 2016.
- [14] E. A. Huerta et al. Complete waveform model for compact binaries on eccentric orbits. *Phys. Rev. D*, 95:024038, Jan 2017.
- [15] E.A. Huerta and D.A. Brown. Effect of eccentricity on binary neutron star searches in Advanced LIGO. *Phys. Rev. D*, 87(12):127501, 2013.
- [16] E. A. Huerta et al. Accurate and efficient waveforms for compact binaries on eccentric orbits. *Phys. Rev. D*, 90(8):084016, 2014.
- [17] D. George and E.A. Huerta. Deep Neural Networks to Enable Real-time Multimessenger Astrophysics. *Phys. Rev. D*, 97(4):044039, 2018.
- [18] D. George and E.A. Huerta. Deep Learning for Real-time Gravitational Wave Detection and Parameter Estimation: Results with Advanced LIGO Data. *Phys. Lett. B*, 778:64–70, 2018.
- [19] K. Fukushima. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics*, 36:193–202, 1980.
- [20] Y. LeCun et al. Object recognition with gradient-based learning. In *Shape, Contour and Grouping in Computer Vision*, pages 319–345, 1999.
- [21] R.E. Colgan et al. Efficient Gravitational-wave Glitch Identification from Environmental Data Through Machine Learning. *Phys. Rev. D*, 101(10):102003, 2020.
- [22] D.E. Rumelhart, G.E. Hinton, and R.J. Williams. *Learning Internal Representations by Error Propagation*, page 318–362. MIT Press, Cambridge, MA, USA, 1986.
- [23] F. Morawski et al. Anomaly detection in gravitational waves data using convolutional autoencoders. *Machine Learning: Science and Technology*, 2021.
- [24] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9:1735–80, 12 1997.
- [25] J. Chung et al. Empirical evaluation of gated recurrent neural networks on sequence modeling. *CoRR*, abs/1412.3555, 2014.
- [26] G. Ashton et al. Bilby: A user-friendly bayesian inference library for gravitational-wave astronomy. *The Astrophysical Journal Supplement Series*, 241(2):27, apr 2019.
- [27] I.M. Romero-Shaw et al. Bayesian inference for compact binary coalescences with bilby: validation and application to the first LIGO–Virgo gravitational-wave transient catalogue. *Monthly Notices of the Royal Astronomical Society*, 499(3):3295–3319, 09 2020.
- [28] Z. Benkő, T. Bábel, and Z. Somogyvári. How to find a unicorn: a novel model-free, unsupervised anomaly detection method for time series. <https://arxiv.org/abs/2004.11468>, 2020. arXiv cs.LG/2004.11468.
- [29] K. Pearson F.R.S. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901.
- [30] J. Powell et al. Classification methods for noise transients in advanced gravitational-wave detectors. *Classical and Quantum Gravity*, 32(21):215012, oct 2015.
- [31] Jade Powell, Alejandro Torres-Forné, Ryan Lynch, Daniele Trifirò, Elena Cuoco, Marco Cavaglia, Ik Siong Heng, and José A Font. Classification methods for noise transients in advanced gravitational-wave detectors II: performance tests on advanced LIGO data. *Classical and Quantum Gravity*, 34(3):034002, jan 2017.

- [32] N.S. Philip and K.B. Joseph. Boosting the differences: A fast bayesian classifier neural network. *Intell. Data Anal.*, 4(6):463–473, December 2000.
- [33] N. Mukund et al. Transient classification in ligo data using difference boosting neural network. *Phys. Rev. D*, 95:104059, May 2017.
- [34] T. Gebhard and N. Kilbertus. Generate Gravitational-Wave Data (GGWD). <https://github.com/timothygebhard/ggwd>, 2019.
- [35] Eric Moreno. Source-Agnostic Gravitational-Wave Detection with Recurrent Autoencoders: H1 detector dataset, July 2021. <https://doi.org/10.5281/zenodo.5121514>.
- [36] Eric Moreno. Source-Agnostic Gravitational-Wave Detection with Recurrent Autoencoders: L1 detector dataset, July 2021. <https://doi.org/10.5281/zenodo.5121510>.
- [37] A. Nitz et al. gwastro/pycbc: Pycbc release v1.16.9, August 2020.
- [38] LIGO Scientific Collaboration. LIGO Algorithm Library - LALSuite. free software (GPL), 2018.
- [39] A. Bohé et al. Improved effective-one-body model of spinning, nonprecessing binary black holes for the era of gravitational-wave astrophysics with advanced detectors. *Phys. Rev. D*, 95(4):044028, 2017.
- [40] B.P. Abbott et al. Properties of the binary neutron star merger GW170817. *Phys. Rev. X*, 9(1):011001, 2019.
- [41] E. Cuoco et al. On line power spectra identification and whitening for the noise in interferometric gravitational wave detectors. *Class. Quant. Grav.*, 18:1727–1752, 2001.
- [42] Zhiqiang Que et al. Accelerating recurrent neural networks for gravitational wave experiments, 2021.
- [43] O. Cerri et al. Variational Autoencoders for New Physics Mining at the Large Hadron Collider. *JHEP*, 05:036, 2019.