


A Deep Transfer Learning Approach on Identifying Glitch Wave-form in Gravitational Wave Data

Reymond Mesuga*  and Brian James Bayanay

Department of Physical Sciences, College of Science

Polytechnic University of the Philippines, Manila, Philippines

*Corresponding author: rmesuga@iskolarngbayan.pup.edu.ph

Abstract—LIGO interferometer is considered the most sensitive and complicated gravitational experimental equipment ever built. Its main objective is to detect the gravitational wave from the strongest events in the universe by observing if the length of its 4-kilometer arms change by a distance 10,000 times smaller than the diameter of a proton. Due to its sensitivity, interferometer is prone to the disturbance of external noises which affects the data being collected to detect the gravitational wave. These noises are commonly called by the gravitational-wave community as glitches. This study focuses on identifying those glitches using different deep transfer learning algorithms. The extensive experiment shows that algorithm with architecture VGG19 recorded the highest AUC-ROC among other experimented algorithm with 0.9898. While all of the experimented algorithm achieved a considerably high AUC-ROC, some of the algorithm suffered from class imbalance of the dataset which has a detrimental effect when identifying other classes.

Index Terms—Transfer Learning, LIGO, Gravitational Wave, Glitch Wave-Forms

I. INTRODUCTION

A. Physics Background

Gravitational waves (GWs) are deformations in spacetime that result from astrophysical phenomena involving celestial objects of masses much heavier than that of the sun moving at speeds up to a significant fraction of the speed of light, mainly called compact objects. GWs result from either mergers of binaries of compact objects, such as binary black hole (BBH) mergers, binary neutron star (BNS) mergers, neutron star-black hole binaries, white dwarf binaries, etc., or from self-production by a massive release of energy from astrophysical phenomena like stellar collapse (i.e., supernovae).

Ever since the first direct detection of GWs by the joint collaboration of LIGO and Virgo on September 14, 2015, the field of gravitational-wave astronomy has become one of the rising fields of research in contemporary physics, and with upgrades to the LIGO detectors in the US as well the VIRGO detector in Italy, combined with the newly-operational KAGRA Observatory in Japan and the operation of the LISA Mission in future years, more and more GW events are being and will be detected, and with these detections come terabytes of data that are in great need of accurate analysis, to ensure that the signals that these observatories are indeed signals from outer space and not noise, either of terrestrial or electromagnetic origin. To build a somewhat good foundation

of how the GW event data is being gathered, it is a must to give a short elaboration on the experimental setup of these observatories.

In the case of LIGO-Hanford and LIGO-Livingston, based in the United States, the 2 observatories are ground-based Michelson interferometers with arms spanning 4 km, where a 20-watt laser is fired, passing through a power recycling mirror, which then fully transmits light incident from the laser and reflects light from the other side increasing the power of the light field between the mirror and the subsequent beam splitter. From the beam splitter, the light travels along two orthogonal arms, and by using partially reflecting mirrors, Fabry-Pérot cavities are created in both arms that increase the effective path length of laser light in the arm. When a GW of sufficient energy passes through the interferometer, the spacetime in the local area is deformed, manifested through the effective change in length of one or both Fabry-Pérot cavities. This change in length will cause the light in the cavity to be slightly out of phase with the incoming light, which will lead to the cavity/s being out of coherence, and the laser light, which are tuned to destructively interfere at the detectors, will have a slightly periodically varying detuning, resulting in a measurable signal, with the detectors' sensitivities up to lengths 10000 times smaller than the diameter of the proton [1]. Due to this and the LIGO-US detectors being ground-based, factors such as instrument noise and environmental influence [2] to name a few, the LIGO detectors not only records the GW strain data, but also over 200,000 auxiliary channels that monitor instrument behavior and environmental conditions [3]. Then, the GW strain data and the data from the auxiliary channels (which may or may not contain legitimate GW strain data) are then subject to data analysis.

In the analysis of GW data, of particular concern are transient, non-Gaussian noise features, called glitches, which are instrumental or environmental in nature (caused by e.g., small ground motions, ringing of the test-mass suspension system at resonant frequencies, or fluctuations in the laser) and come in a wide variety of time-frequency amplitude morphologies [4], and can mimic true gravitational wave signals and can hinder sensitivity conditions [5]. These glitches are then classified by common origin and/or similar morphological characteristics [4].

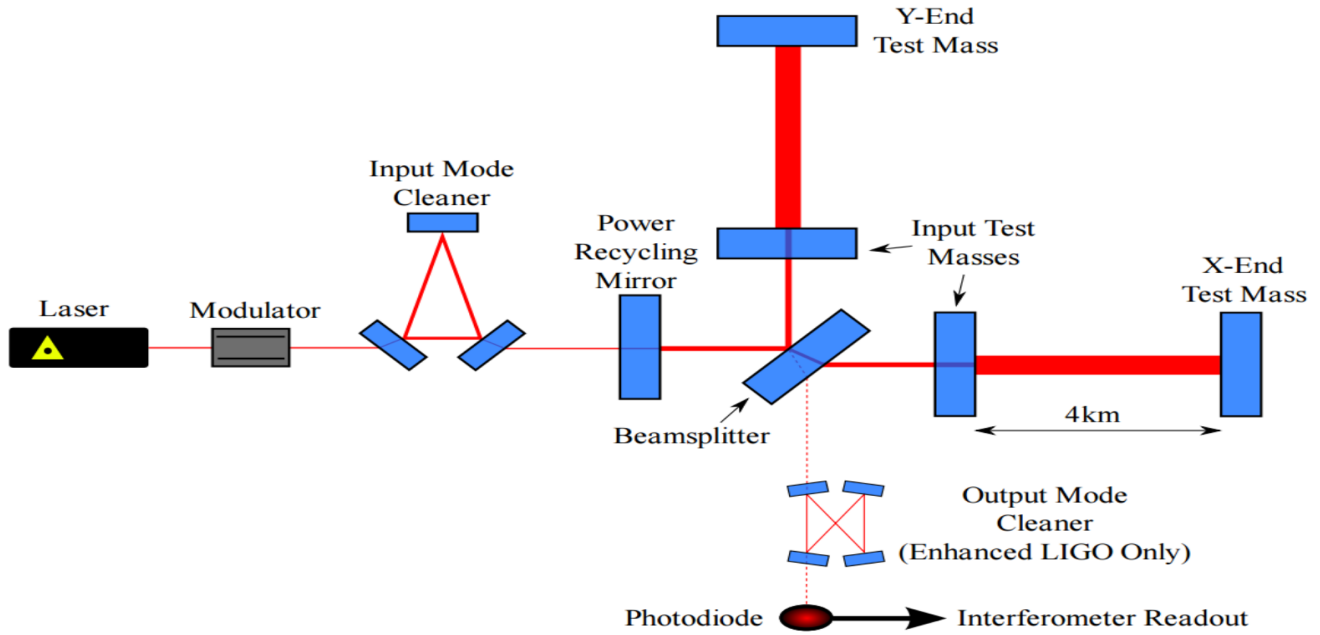


Figure 1. Experimental setup of the Advanced LIGO Detector, [1]

B. Deep Learning

The most used method used in the identification and classification of glitches in GW data is by means of machine learning (ML) algorithms [5]–[11] such as dictionary learning [5], similarity learning [6], deep transfer learning [7] among many a method. Building on this, the LIGO-Virgo Collaboration has advocated citizen science involvements in the search for these glitches, e.g., Gravity Spy [1], which speeds up the process of refining the increasingly large amount of GW data.




Figure 2. Comparison of images from a glitch signal (left) and a GW signal (right) in LIGO O2 data, [26]

The refining of GW data is of paramount importance to the scientific community, particularly to physicists working in general relativity, cosmology, astrophysics, quantum gravity, etc. because the phenomenon itself serves as a validity of theories of gravity, mainly general relativity, and opens new questions to various fields in physics, like the mass threshold at which GW should occur, one of particular interest is the most recent detection (O3) of LIGO-Virgo, GW190814 [12], which indicates a “mass gap” between the heaviest neutron stars and the lightest black holes. These GW data is also used

in the investigation of the Hubble expansion of the universe [13], [14], cosmic inflation via the existence of a stochastic gravitational-wave background [15], [16], existence of dark matter [17], among many new physics. With that being said, the efficiency of the different machine learning algorithms used in searching, classifying, and mitigating glitches should be considered, in order to obtain the most accurate data possible, and for this purpose, this experiment aims to compare some of the different machine learning algorithms used in noise detection in GW data, and their respective efficiencies in doing so.

The use of deep learning algorithms will be adopted in this paper. Deep learning (DL) is a type of machine learning (ML) algorithm where it uses a so called artificial neural network (ANN) to learn from different input data (i.e., images, sounds and texts) [18], [19]. ANN was inspired on how the human brain works. In fact, ANN has its own version of neuron that functions almost similar to biological neuron called artificial neuron which represents the nodes that can be found in the hidden layer (see Fig. 3). In a biological neuron, if the signals of information received by the synapse are strong enough (or surpass a certain threshold) [19]. When it comes to artificial neuron, if the weight of an input is not enough then the neuron will not be activated. In an image classification problem, each pixels on the image will serve as the inputs in input layer (see Fig. 3).

After setting the pixels of an image as inputs, the ANN will then apply a randomly selected weights and multiply each on inputs. This can be defined as:

$$y = w * x + b \quad (1)$$




Figure 3. A simple diagram of artificial neural network, [27]

where y is the output, w is the weight, x is the input and b is a special kind of weight called bias. This function is also the equation of the line.

Unfortunately, ANN (or/and DL algorithms in general) won't give any meaningful results just by just using a linear transformation. The application of non-linearity is needed because inputs such as images, sounds and texts are naturally non-linear. To apply non-linearity, each neuron on the hidden layer needs a so called activation function. Now, the output y on each neuron can be defined as:

$$y = \varphi(x * w + b) \quad (2)$$

where φ is the activation function. Note that equation 2 only represents an output y using a single input x . In general, output y of a neuron is the summation of all input x with their corresponding weights and bias and can be defined as equation 3:

$$y_k = \varphi\left(\sum_{i=0}^m x_i * w_{ki} + b_k\right) \quad (3)$$

$$\sum_{i=0}^m x_i * w_{ki} + b_k \quad (4)$$

The most common and up-to-date activation function is the ReLU activation function. This can be defined as:

$$f(x) = \max(0, x) \quad (5)$$

The idea here is that, if the value of equation 4 is less than or equal to zero, the output y will be automatically set as zero and will be deactivated. Otherwise, if it is greater than zero, then the output y will stay as it is and hence activated [20].

Now, one cannot expect a deep learning algorithm to have an accurate and meaningful result just by using a random

selected weights. In order for a deep learning algorithm to have a more accurate result, the algorithm needs to be trained. Training the algorithm involves the adjustments of the weights. These weights need to be adjusted in a way that it fits on what is needed by the algorithm to give a more accurate result. These adjustments will continue as long as the algorithm has not reached yet its minimum loss. The measurement of loss will be done using a *Loss Function* which measures how good or bad a deep learning is to classify each class on the dataset. The most common and up-to-date loss function used in a multi-class image classification is the *Categorical Cross Entropy Loss* [22]. Meanwhile, the most common and up-to-date optimization algorithm that is used to adjust the weights is called the *Adam Optimizer* [21]. This is the simple approach to understand how artificial neural network (ANN) works in an image classification. That is, by setting the pixels of an image as inputs, applying random weights on each input, applying non-linear transformation on each input with their corresponding weights using an activation function, training the model to lower the loss measured by a loss function and adjusting the weights using the optimizer.




Figure 4. A simple diagram of convolutional neural network (CNN), [23]

Artificial neural network (ANN) has many variations of algorithms. The most successful ANN algorithm is called Convolutional Neural Network (also called as CNN or ConvNet) [23], [24]. The structure of CNN can be divided into two parts which are the base and the head. The base of CNN is used to extract the features from an image and is formed primarily of three basic components namely *convolution layer*, *ReLU activation function* and *maximum pooling*. On the other hand, the head of the CNN is responsible for determining the class of the image. The main usage of the convolution layer is to filter an image for a particular feature. Meanwhile, ReLU activation detects the feature within the filtered image and maximum pooling is responsible for the enhancement [24].

The general objective of the study is to evaluate the performance of various deep transfer learning models in glitch waveform detection in gravitational-wave data. The specific objectives are as follows: a) Evaluate and Compare the performance of different algorithms using specific evaluation metrics, b) Identify how the quantity of the data affects the performance of the algorithm, and c) Identify how the complexity of the

algorithm affects its performance in identifying glitch waveforms.

II. METHODOLOGY

A. Data Gathering and Preparation

The dataset used in this study were gathered from a Kaggle repository that have been classified as a part of the Gravitational Spy Dataset [28]. The only difference between the Kaggle and original version of dataset is that image data found in a Kaggle repository has no axes and was already divided into train set, validation set and test set which contains 22348 (70%), 4800 (15%), and 4720 (15%) respectively with the total of 31868 image data. Table I shows the number of data per classes. As observed, the class Blip contains the majority of images which has 1821 images and the classes 1400Ripples, None_of_the_Above, Chirp, Air_Compressor, Wandering_Line and Paired_Doves did not even make it above a hundred.

Table I
NO. OF DATA PER CLASSES

Glitch Classes	No. of Data
Blip	1821
Koi_Fish	706
Low_Frequency_Burst	621
Light_Modulation	512
Power_Line	449
Extremely_Loud	447
Low_Frequency_Lines	447
Scattered_Light	443
Violin_Mode	412
Scratchy	337
1080Lines	328
Whistle	299
Helix	279
Repeating_Blips	263
No_Glitch	150
Tomte	103
None_of_the_Above	81
1400Ripples	81
Chirp	60
Air_Compressor	58
Wandering_Line	42
Paired_Doves	27

B. Deep Transfer Learning Algorithms

The type of deep learning algorithms that will be used in the experiment are all deep transfer learning algorithm. Deep transfer learning or simply transfer learning uses pre-trained architectures as its base algorithm. As mentioned in section I-B, the structure or architecture of convolutional neural network (CNN) can be divided into two parts (i.e., base and head). Most transfer learning algorithm uses CNN architecture as well. The only difference is that, the base of transfer learning algorithms were trained already using different images in the past. The most notable dataset that is commonly used to train a base for a pre-trained algorithm is called ImageNet which contains 1.2 million images that has 1000 different classes [25]. The architectures tested were ResNet [29], VGG [30], Xception [31], Inception [32], and DenseNet [33]. To

be specific, the following pre-trained architectures used are as follows: ResNet101, ResNet101V2, ResNet152, ResNet50, ResNet50v2, VGG16, VGG19, Xception, InceptionResnetV2, and DenseNet169.

C. Performance Metrics

The below figures (i.e., Fig. 5 and 6) shows the illustration for confusion matrix and AUC-ROC will be discussed further in the following sub-subsection.

		PREDICTED classification			
		a	b	c	d
ACTUAL classification	Classes	a	b	c	d
	a	TN	FP	TN	TN
	b	FN	TP	FN	FN
	c	TN	FP	TN	TN
d	TN	FP	TN	TN	

Figure 5. Confusion matrix for multi-classes, [34]




Figure 6. AUC-ROC, [35]

1) **Confusion Matrix:** Confusion matrix for multi-class classification works the same way as in binary classification. It can be visualize in Fig. 5 where y -axis corresponds to the actual classification of samples while x -axis corresponds to the predicted classification.

The main goal of using a confusion matrix is to visualize the quantity of samples that are correctly and wrongly predicted

by the classifier. The samples in confusion matrix can be categorized into four types namely *True Positive* (TP), *True Negatives* (TN), *False Positives* (FP), and *False Negatives* (FN).

2) **Area Under the Curve of Receiver Operating Characteristic (AUC-ROC):** AUC-ROC is primarily used in binary classification problems classification according to their one-vs-all precision-recall-curves. AUC-ROC is primarily used to see if the classifier was able to separate the classes in the data. In order to plot the actual ROC, *True Positive Rate* (TPR) and *False Positive Rate* (FPR) must be computed with many different threshold and plotted as in Fig. 6 where x - axis corresponds the the FPR while y - axis corresponds to TPR. The ideal value of AUC-ROC is 1 which means that the classifier was able to separate perfectly the classes in the data. The lowest is 0.5 which means the classifier is just making predictions randomly (see Figure 6. The AUC-ROC metric was used as the primary performance metric in this study as it is commonly used in field of deep learning and is better than using accuracy for decision making.

III. DISCUSSION OF RESULTS

The bar plot in Fig. 7 shows that VGG19 recorded the highest AUC-ROC overall followed by ResNet101V2, ReNet50V2, and DenseNet169 all of which achieved an AUC-ROC higher than 0.98. Meanwhile, ResNet152 recoded the lowest AUC-ROC of 0.93954 which is not bad. Overall, all the algorithm have recorded an AUC-ROC higher than 0.90 which is considerably high or decent.

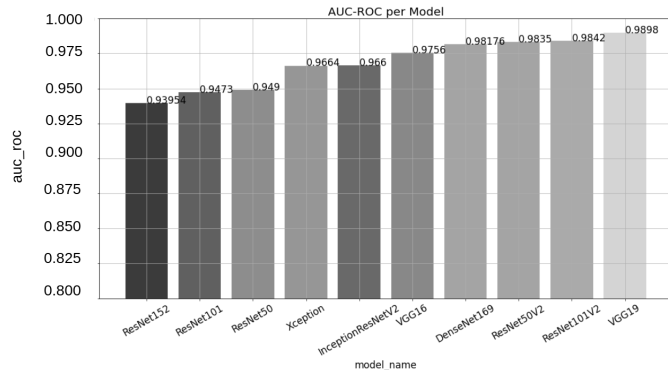


Figure 7. AUC-ROC per algorithm

To better understand the difference in performance between each algorithm, it is also good to plot their AUC-ROC vs Complexity. The term complexity here means the total parameters of the model which can be interpreted as how complex they are compared to others.

As shown in Fig. 8, lighter algorithm architectures like VGG16 and VGG19 outperformed several architectures like InceptionResNetV2, Xception, ResNet50, ResNet101, and ResNet52 despite having a more complex or bigger network.

Fig. 9 shows the confusion matrix of VGG19. It is visible that it is capable of identifying low quantity classes such as the class *Air_Compressor* and *Wandering_Line*. Note that




Figure 8. AUC-ROC vs Complexity

image data in Gravity Spy Dataset has a problem in class imbalanced (see Table I). The two mentioned classes are one of the classes with lowest number of data which greatly affect the performance of other algorithms. Despite that, VGG19 still able to identify those classes at considerably decent performance.




Figure 9. Confusion matrix of VGG19 (AUC-ROC=0.9898)

Meanwhile, Fig. 10 shows the confusion matrix of ResNet152 which achieved an AUC-ROC of 0.93954 which is the lowest recorded AUC-ROC among other experimental algorithms. It was visible that ResNet152 were not able to properly almost half of the classes in the Gravity Spy Dataset. Despite that, it still achieved a considerably high AUC-ROC due to the fact that the upper half classes in the confusion matrix contains majority of the samples in the dataset.

With that, it could be interpreted that in order to use deep transfer learning algorithms in a dataset that contains many classes with imbalanced samples, AUC-ROC have to be higher as much as possible with, at least, 0.97 or above in order to perform decently in the real world.

VII. CONFLICT OF 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.

REFERENCES

- [1] "Zooniverse," Zooniverse.org, 2021. <https://www.zooniverse.org/projects/zooniverse/gravity-spy/about/research> (accessed Jul. 04, 2021).
- [2] S. Doughton, "Ripples in space-time or 3-pound bird? Ravens at Hanford foul test of Einstein's theory," *The Seattle Times*, May 14, 2018. <https://www.seattletimes.com/seattle-news/science/suddenly-there-came-a-tapping-ravens-cause-blipsin-massive-physics-instrument-at-hanford/> (accessed Jul. 04, 2021).
- [3] B. P. Abbott et al., "Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914," *Classical and Quantum Gravity*, vol. 33, no. 13, p. 134001, Jun. 2016, doi: 10.1088/0264-9381/33/13/134001
- [4] M. Zevin, S. Coughlin, S. Bahaadini, E. Besler, N. Rohani, S. Allen, M. Cabero, K. Crowston, A. K. Katsaggelos, S. L. Larson, T. K. Lee, C. Lintott, T. B. Littenberg, A. Lundgren, C. Østerlund, J. R. Smith, L. Trouille, and V. Kalogera, "Gravity Spy: integrating advanced LIGO detector characterization, machine learning, and citizen science," *Classical and Quantum Gravity*, vol. 34, no. 6, p. 064003, 2017.
- [5] M. Llorens-Monteagudo, A. Torres-Forné, J. A. Font, and A. Marquina, "Classification of gravitational-wave glitches via dictionary learning," *Classical and Quantum Gravity*, vol. 36, no. 7, p. 075005, 2019.
- [6] S. Coughlin, S. Bahaadini, N. Rohani, M. Zevin, O. Patane, M. Harandi, C. Jackson, V. Noroozi, S. Allen, J. Areeda, M. Coughlin, P. Ruiz, C. P. L. Berry, K. Crowston, A. K. Katsaggelos, A. Lundgren, C. Østerlund, J. R. Smith, L. Trouille, and V. Kalogera, "Classifying the unknown: Discovering novel gravitational-wave detector glitches using similarity learning," *Physical Review D*, vol. 99, no. 8, 2019
- [7] George, Daniel, Shen, Hongyu, Huerta, E.A. (2017). Glitch Classification and Clustering for LIGO with Deep Transfer Learning. Workshop on Deep Learning for Physical Sciences (DLPS 2017), NIPS 2017, Long Beach, CA, USA.
- [8] R. E. Colgan, K. R. Corley, Y. Lau, I. Bartos, J. N. Wright, Z. Márka, and S. Márka, "Efficient gravitational-wave glitch identification from environmental data through machine learning," *Physical Review D*, vol. 101, no. 10, 2020.
- [9] D. Davis, L. V. White, and P. R. Saulson, "Utilizing aLIGO glitch classifications to validate gravitational-wave candidates," *Classical and Quantum Gravity*, vol. 37, no. 14, p. 145001, 2020.
- [10] S. Bahaadini, N. Rohani, S. Coughlin, M. Zevin, V. Kalogera, and A. K. Katsaggelos, "Deep multiview models for glitch classification," 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2017.
- [11] S. Bahaadini et al., "Machine learning for Gravity Spy: Glitch classification and dataset," *Information Sciences*, vol. 444, pp. 172–186, May 2018, doi: 10.1016/j.ins.2018.02.068.
- [12] R. Abbott et al., "GW190814: Gravitational Waves from the Coalescence of a 23 Solar Mass Black Hole with a 2.6 Solar Mass Compact Object," *The Astrophysical Journal*, vol. 896, no. 2, p. L44, Jun. 2020, doi: 10.3847/2041-8213/ab960f.
- [13] J. Calderón Bustillo, S. Leong, T. Dietrich and P. Lasky, "Mapping the Universe Expansion: Enabling Percent-level Measurements of the Hubble Constant with a Single Binary Neutron-star Merger Detection", *The Astrophysical Journal Letters*, vol. 912, no. 1, p. L10, 2021. Available: 10.3847/2041-8213/abf502.
- [14] W. M. Farr, M. Fishbach, J. Ye, and D. E. Holz, "A Future Percent-level Measurement of the Hubble Expansion at Redshift 0.8 with Advanced LIGO," *The Astrophysical Journal*, vol. 883, no. 2, p. L42, Oct. 2019, doi: 10.3847/2041-8213/ab4284.
- [15] "Gravitational waves from inflation," *Www.sif.it*, 2015. <https://www.sif.it/riviste/sif/ncr/econtents/2016/039/09/article/0> (accessed Jul. 04, 2021).
- [16] J. R. H. Brandenberger, "Is the spectrum of gravitational waves the 'Holy Grail' of inflation?," *The European Physical Journal C*, vol. 79, no. 5, May 2019, doi: 10.1140/epjc/s10052-019-6883-4.




Figure 10. Confusion matrix of ResNet152 (AUC-ROC=0.93954)

IV. CONCLUSION

There are still room for improvements when it comes to the domain of deep transfer learning applications in identifying glitch wave-form in gravitational wave data. The need for improvement is needed especially on the part of producing more data to train algorithms. Using a decent quantity of data with balanced classes is really a necessity and will definitely help the performance of the algorithm as it is one of the major problems found in this study. All of the algorithms featured in this study shows a promising performance that could potentially useful for the use of gravitational-wave scientists. But, there are still differences in performances between algorithms in which some perform better than the other. The differences on the performances of each algorithm should one of the major factor when applying deep transfer learning in production specifically on identifying glitch wave-forms in gravitational wave data.

V. ACKNOWLEDGEMENT

The authors would like to thank Mr. Mark Anthony Burgonio, MS for allowing this study to be conducted under his Advance Laboratory course.

VI. CREDIT AUTHORSHIP CONTRIBUTION STATEMENT

Reymond Mesuga: Conceptualization, Methodology, Software, Investigation, Validation, Formal Analysis, Resources, Data Curation, Writing-Original Draft, Writing-Review & Editing, Visualization, Project Administration, Supervision; **Brian James Bayanay** Investigation, Writing-Original Draft, Writing-Review and Editing;

-
- [17] Y. Michimura, T. Fujita, S. Morisaki, H. Nakatsuka, and I. Obata, "Ultralight vector dark matter search with auxiliary length channels of gravitational wave detectors," *Physical Review D*, vol. 102, no. 10, Nov. 2020, doi: 10.1103/physrevd.102.102001.
- [18] A. U. Ruby, D. I. Prasannavenkatesan Theerthagiri, and Y. Vamsidhar, Binary cross entropy with deep learning technique for image classification, *International Journal* 9 (2020).
- [19] M. C. Nwadiugwu, "Neural Networks, Artificial Intelligence and the Computational Brain," arXiv.org, 2020. <https://arxiv.org/abs/2101.08635> (accessed Jul. 04, 2021).
- [20] V. Nair and G. Hinton, "Rectified Linear Units Improve Restricted Boltzmann Machines,." [Online]. Available: <https://www.cs.toronto.edu/~fritz/absps/reluICML.pdf>.
- [21] I. Kandel, M. Castelli, and A. Popovič, Comparative study of first order optimizers for image classification using convolutional neural networks on histopathology images, *Journal of Imaging* 6, 92 (2020).
- [22] A. U. Ruby, D. I. Prasannavenkatesan Theerthagiri, and Y. Vamsidhar, Binary cross entropy with deep learning technique for image classification, *International Journal* 9 (2020).
- [23] V. Phung and E. Rhee, A High-Accuracy Model Average Ensemble of Convolutional Neural Networks for Classification of Cloud Image Patches on Small Datasets, *Applied Sciences* 9 (2019)
- [24] A. Ghosh, Fundamental Concepts of Convolutional Neural Network, Recent Trends and Advances in Artificial Intelligence and Internet of Things, (2020)
- [25] A. Krizhevsky, ImageNet Classification with Deep Convolutional Neural Networks, *NeurIPS Proceedings*, (2012)
- [26] R. E. Colgan et al., "Efficient gravitational-wave glitch identification from environmental data through machine learning," *Physical Review D*, vol. 101, no. 10, May 2020, doi: 10.1103/physrevd.101.102003.
- [27] M. H. Esfe, S. A. Eftekhari, M. Hekmatifar, and D. Toghraie, "A well-trained artificial neural network for predicting the rheological behavior of MWCNT–Al₂O₃ (30–70%)/oil SAE40 hybrid nanofluid," *Scientific Reports*, vol. 11, no. 1, p. 17696, Aug. 2021, doi: 10.1038/s41598-021-96808-4.
- [28] Coughlin, Scott. (2018). Updated Gravity Spy Data Set (v1.1.0) [Data set]. Zenodo. <https://doi.org/10.5281/zenodo.1476551>
- [29] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," Dec. 2015. [Online]. Available: <https://arxiv.org/pdf/1512.03385.pdf>.
- [30] K. Simonyan and A. Zisserman, "Published as a conference paper at ICLR 2015 VERY DEEP CONVOLUTIONAL NETWORKS FOR LARGE-SCALE IMAGE RECOGNITION," 2015. [Online]. Available: <https://arxiv.org/pdf/1409.1556.pdf>.
- [31] F. Chollet, "Xception: Deep Learning with Depthwise Separable Convolutions." [Online]. Available: <https://arxiv.org/pdf/1610.02357.pdf>.
- [32] C. Szegedy et al., "Going deeper with convolutions," 2014. [Online]. Available: <https://arxiv.org/pdf/1409.4842.pdf>.
- [33] G. Huang, Z. Liu, L. Van Der Maaten, and K. Weinberger, "Densely Connected Convolutional Networks," 2018. [Online]. Available: <https://arxiv.org/pdf/1608.06993.pdf>.
- [34] M.-C. Metrics For and Classification, AN OVERVIEW, 2020. [Online]. Available: <https://arxiv.org/pdf/2008.05756.pdf>.
- [35] Wikipedia Contributors, "File:Roc-draft-xkcd-style.svg," Wikipedia, 2022. <https://en.wikipedia.org/wiki/File:Roc-draft-xkcd-style.svg> (accessed Apr. 30, 2022).