



Deep Learning for Cardiologist-Level Myocardial Infarction Detection in Electrocardiograms

Arjun Gupta^{1,2(✉)}, Eliu Huerta^{1,3,4}, Zhizhen Zhao^{1,2}, and Issam Moussa^{1,5,6}

¹ NCSA Center for Artificial Intelligence Innovation,
University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA
arjung2@illinois.edu

² Department of Electrical and Computer Engineering,
University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

³ Department of Astronomy, University of Illinois at Urbana-Champaign,
Urbana, IL 61801, USA

⁴ Department of Physics, University of Illinois at Urbana-Champaign,
Urbana, IL 61801, USA

⁵ Beckman Institute for Advanced Science and Technology,
University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

⁶ Carle Illinois College of Medicine, University of Illinois at Urbana-Champaign,
Urbana, IL 61820, USA

Abstract. Myocardial infarction is the leading cause of death worldwide. In this paper, we design domain-inspired neural network models to detect myocardial infarction. First, we study the contribution of various leads. This systematic analysis, first of its kind in the literature, indicates that out of 15 ECG leads, data from the v6, vz, and ii leads are critical to correctly identify myocardial infarction. Second, we use this finding and adapt the ConvNetQuake neural network model—originally designed to identify earthquakes—to attain state-of-the-art classification results for myocardial infarction, achieving 99.43% classification accuracy on a record-wise split, and 97.83% classification accuracy on a patient-wise split. These two results represent cardiologist-level performance level for myocardial infarction detection after feeding only 10 s of raw ECG data into our model. Third, we show that our multi-ECG-channel neural network achieves cardiologist-level performance without the need of any kind of manual feature extraction or data pre-processing.

Keywords: Machine learning · Signal processing · Biomedical engineering

1 Introduction

Heart disease is the leading cause of death worldwide. Amongst patients with cardiovascular diseases, myocardial infarction is the main cause of death. In

order to provide adequate healthcare support to patients who may experience this clinical event, it is essential to gather supportive evidence in a timely manner to help secure a correct diagnosis.

Deep learning is the engine of the artificial intelligence renaissance we have experienced since the early 2010s [21,31]. Deep learning applications are ubiquitous in modern technologies, and there is a sense of urgency to better understand how these algorithms may be designed, trained and utilized in an optimal manner—an emergent area of research dubbed “scientific machine learning” [38].

The confluence of mathematical and statistical models, along with numerical simulations to inform the design and training of deep learning algorithms is a booming enterprise [16,25,26,34]. Furthermore, the development of novel schemes to accelerate the training of deep learning models using hundreds to thousands of graphics processing units (GPUs) in high performance computing platforms has enabled the construction of more sophisticated neural network models, trained with TB-size data sets, that become more resilient to noise anomalies, and which exhibit state-of-the-art results for classification and regression analyses [5,13,15,19]. An additional advantage of these algorithms is that once the compute-intensive training stage is complete, they require minimal computational resources for inference analyses [4,6–8,8,9,11,12,17,28,33,35,39,40].

The combination of data fusion and machine learning exhibit great promise to enable innovation in healthcare. However, in order to realize this expectation, it is essential to first ascertain what AI architectures and data sets provide the most critical information to address a given challenge. In the case of diagnosis of myocardial injury, we require a number of supportive evidence in terms of typical symptoms, electrocardiographic (ECG) changes, imaging evidence that indicates new loss of viable myocardium or new regional wall motion abnormality, etc. Along this line of thought, in this article we:

1. Fill a void in the literature regarding the use of ECG time-series data to identify heart conditions, i.e., in our literature review, we were unable to find a convincing line of argumentation regarding the choice of ECG data used in previous machine learning studies to identify heart conditions. Given that the choices in previous analyses are arbitrary, we present a detailed analysis of the contribution of each ECG lead for the identification of heart conditions.
2. Illustrate the benefits of transferring knowledge between disparate areas of research that are threaded by a common theme, namely, the use of multiple data channels of information to enhance our confidence in the prediction of deep learning models. In this particular case, we adapt a neural network model that was designed to identify earthquakes by feeding into the model multiple channels of information. In our case, we first identify the ECG leads that provide the highest confidence to tell apart healthy from unhealthy hearts, and then re-design the neural network model to process simultaneously the top three ECG leads that provide the highest classification results.
3. Show that feature engineering is not necessary to curate the ECG data used to train, validate and test our neural network models, which currently represent cardiologist state-of-the-art performance for myocardial infarction detection.

2 Related Work

Myocardial infarction detection on the PTB database is a problem that has been studied extensively in the past. [1] were one of the first to develop a convolutional neural network for the automatic detection of myocardial infarction. This was done on both noisy and de-noised electrocardiogram signals, without any manual feature extraction. [30] used two features-the T-wave integral and the total integral-for the localization and detection of myocardial infarction. Using these two features and various statistical learning techniques such as artificial neural networks, probabilistic neural networks, k-nearest neighbors, multi-layer perceptrons, and naive-Bayes classifiers, [30] were able to achieve a 74% localization accuracy and a 94% detection accuracy. [20] developed two different kinds of artificial neural networks-a radial basis function (RBF) and a multi-layer perceptron (MLP)-for the detection of myocardial infarction. [36] argue that supervised learning techniques have achieved limited success and apply multiple instance learning (MIL) to automatic electrocardiogram classification. They demonstrate that they're proposed algorithm called LTMIL surpasses previous supervised learning approaches in terms of classification quality.

[22] propose a novel electrocardiogram feature which they obtain by best 20th-order polynomial approximation of a given ECG signal. Developing a classification model using this feature, [22] are able to achieve a 94.4% accuracy for myocardial infarction detection. [32] use stationary wavelet transform to decompose a given electrocardiogram signal into different sub-bands. Then, support vector machine (SVM) and k-nearest neighbor models are developed to classify the presence of myocardial infarction using sample entropy, normalized sub-band energy, log energy entropy, and median slope from the selected bands as features. [18] use transfer learning from a model trained for arrhythmia classification to develop a model capable of detecting myocardial infarction. A deep convolutional neural network model is used, achieving a 95.9% classification accuracy for myocardial infarction detection.

[29] use a simple adaptive threshold (SAT) to classify myocardial infarction. A multiresolution approach along with an adaptive thresholding is used to extract the depth of the Q peak and elevation in the ST segment which are the features used for classification. [27] develop a subject-oriented approach based on a convolutional neural network which takes in leads II, III, and AVF. [41] model electrocardiogram signals using second order ordinary differential equations (ODE) and feed the best-fitting time varying coefficients of the ECG signal into a support vector machine (SVM) model. [10] develop a multi-channel convolutional-recurrent neural network model consisting of 16 convolutional layers followed by long-short term memory (LSTM) units to classify the presence myocardial infarction in electrocardiograms. Finally, [23] propose a deep convolutional neural network which takes only 3-s of lead II as input at a time for myocardial infarction detection in ECG signals.

Many approaches [10, 18, 20, 22, 27, 29, 30, 32, 36, 41] extract hand-crafted features from the raw ECG signals before feeding these features into a signal-processing algorithm. Such manual feature extraction is both time-consuming

and not scalable. Our approach differs from others in the community in that it does not require any kind of manual feature extraction. Instead, we adopt a data-driven discovery approach in which we fully exploit the proven capabilities of deep learning algorithms to identify novel features or patterns in data that escape human notice or domain expertise, and which enable the use of neural network models to process data at scale.

On the other hand, the few studies that do not do any feature engineering as a pre-processing step [1,23] do not provide any rationale for their choice of ECG leads to identify or classify heart conditions. To shed new information on this important decision making process, we quantify the contribution of each ECG and Frank lead to correctly identify heart conditions in ECG data. Our results indicate which lead(s) contain the most meaningful information for the detection of myocardial infarction. Based on these findings we design a neural network model that employs these leads to achieve state-of-the-art results.

3 Methods

We begin this section by describing the data curation methodology followed to prepare the training, validation and testing data sets. The ECG time-series data for this analysis was obtained from the public Physikalisch-Technische Bundesanstalt (PTB) database. Thereafter, we describe the architecture of the neural network model used to identify heart conditions. We provide all the key information for reproducibility purposes. All the code developed for this project is open source [3].

3.1 Data Curation

As mentioned before, we use the PTB database, which consists of 549 ECG records from 290 unique patients, with a mean length of over 100 s for each record [14]. The dataset provides data from the 12 conventional ECG leads, along with 3 Frank leads, all sampled at 1000 Hz. Table 1 shows the distribution of heart conditions for the 290 patients in the dataset (Note: the diagnostic class was not available for 22 subjects; these subjects were not used in training.). This data set was used to train, validate and test our neural network models.


During the training stage, a 10-s long two-channel input was fed into the neural network. Both channels were normalized to ensure that the two channels were weighted equally, and time invariance was incorporated by selecting the 10 s long segment randomly from the entire signal. Figure 1 shows sample inputs of the v6 and vz leads for patients with and without heart conditions.

3.2 Neural Network Model

An aspiration in deep learning research is to develop commodity software that may be utilized across disciplines that share common computational grand challenges. In this respect, deep transfer learning has had a number of successful applications [2,13,19,37,40].

Table 1. The diagnostic classes for unique patients in the dataset

Diagnostic class	Number of subjects
Myocardial infarction	148
Cardiomyopathy/Heart failure	18
Bundle branch block	15
Dysrhythmia	14
Myocardial hypertrophy	7
Valvular heart disease	6
Myocarditis	4
Miscellaneous	4
Healthy controls	52

**Fig. 1.** Sample input data for (a) and (c) a healthy patient, and (b) and (d) for an unhealthy patient

Our analysis explores the adequacy of this paradigm by adapting a neural network model that was originally designed to identify earthquakes by processing multiples channels of input data [24]. Herein, we fine-tune the ConvNetQuake,

using an eight-layer convolutional neural network architecture, shown in Fig. 2, which processes 10-s long ECG signals to detect myocardial infarction.




Fig. 2. Architecture of the neural network model used to identify heart conditions feeding 10 s of raw ECG data from the v6 and vz ECG leads. Notice that these two independent time-series data sets are marked by the red and blue input data sets at the top of the diagram. The result is an activation of one of two classes, shown in yellow.

The primary difference between our architecture and that of ConvNetQuake is the use of a one-dimensional batch normalization layer after each convolutional layer to combat over-fitting. While there is much debate about whether batch normalization should proceed or follow the activation function, we observe that for our case, applying batch normalization after the activation yields better results. Another difference from ConvNetQuake is that we employ label smoothing. Label smoothing refers to the act of relaxing the confidence on the labels and is known to help discourage the model from making over-confident predictions. Our experiments showed that both of these techniques helped increase the accuracy of our models.

For training, we used a batch size of 10 and a learning rate of 10^{-4} . The weights were randomly initialized in all cases, and the ADAM optimizer was used. The two classes—myocardial infarction and healthy—are unbalanced in the

database. To overcome this, we sampled input data such that the neural network was exposed to an equal amount of samples from each class. An 80-10-10 train-validation-test split was used in all instances. The analysis was carried out using NVIDIA V100 GPUs, with each experiment taking a few hours to train when trained on one GPU.

It is known that random initialization of the weights at each training round may produce some variance in the performance of neural network models. In view of this observation, we have trained our models from scratch over 100 times to provide a fair representation of the performance of our models. In the results section, we provide average, mean and variance results of these training sweeps. Furthermore, to facilitate reproducibility of our results, we provide all the code used for these studies [3].

4 Results

Here, we evaluate the performance of our models in myocardial infarction classification using standard statistical measures that others have used in the past: the accuracy, sensitivity, specificity, and recall of the models.

4.1 Record-Wise Split

The first stage of our analysis consists of quantifying the contribution of each of the 15 leads—first individually, and then in pairs—to see which lead contains the most meaningful information for the detection of myocardial infarction.

Tables 2 and 3 show our results obtained over a 20-fold cross-validation.

Table 2. Quantification of accuracies for single channels [i - v2] on PTBDB

Lead	i	ii	iii	avl	avr	avf	v1	v2
Average (%)	82.68	89.32	84.91	84.84	86.96	87.44	87.15	83.93
Std. Dev. (%)	6.74	5.52	6.48	3.12	5.23	4.75	4.55	6.38
Median (%)	83.99	90.57	86.56	85.11	88.40	86.94	86.95	82.70

Table 3. As Table 2, for channels [v3 - vz]

Lead	v3	v4	v5	v6	vx	vy	vz
Average (%)	83.40	84.27	91.13	90.84	89.64	87.08	89.41
Std. Dev. (%)	4.58	5.37	6.30	4.68	4.83	5.54	4.49
Median (%)	82.38	84.55	91.84	92.04	89.28	88.38	90.33

Based on these results, we observed that the five channels with the most valuable information for myocardial infarction detection were: v5, v6, vx, vz, and ii.

Next, we formed pairs of channels using the five best channels according to the PTB database and retrained the network with these pairs as input. Our results, presented in Tables 4 and 5 are based on a 20-fold cross validation.

Table 4. Quantification of accuracies for pairs of channels

Leads	v5, v6	v5, vx	v5, vz	v5, ii	v6, vx
Accuracy (%)	92.26	93.77	93.46	94.19	92.98
Std. Dev. (%)	4.78	4.30	4.20	4.85	3.63
Median (%)	92.19	93.51	93.545	93.545	92.74

Table 5. As Table 4, continued

Leads	v6, vz	v6, ii	vx, vz	vx, ii	vz, ii
Accuracy (%)	94.76	93.58	93.44	93.35	92.17
Std. Dev. (%)	3.87	5.09	3.99	4.63	4.23
Median (%)	94.57	93.01	93.42	93.58	91.62

Tables 4 and 5 indicate that when lead v6 and lead vz are paired and fed into the neural network as a 2-channel input, the model is most successful at the task at hand. To see get a more precise estimate of the statistics, we trained a model on the best performing pair of channels 100 times (100-fold cross-validation). Table 6 summarizes our findings.

Table 6. Results based on a 100-fold cross-validation

Statistics	v6, vz
Average (%)	94.67
Std. Dev. (%)	4.04
Median (%)	94.67

As evident, the standard deviation across the 100-fold cross-validation is non-trivial. To figure out whether the standard deviation result was caused by a few outliers, which we confirmed to be the case, we calculated the statistics of the top 20 and top 50 performing models. We present these results in Table 7.

As evident from Table 7, the standard deviations dropped considerably relative to the standard deviation across all 100 models (see Table 6). The much

Table 7. Results based on best performing models

Statistics	Top-50	Top-20
Average (%)	97.91	99.43
Std. Dev. (%)	1.60	0.47
Median (%)	98.04	99.50

lower standard deviation provides evidence for consistent strong performance; many models performed very well, but the overall statistics across 100 models were dragged down by a few models which performed very poorly.

In all of our cross-validation experiments, there are two factors varying across each iteration: the random initialization of the weights of the neural network, and the random train-val-test split. To investigate whether some models were performing poorly due to a poor train-val-test split¹ as opposed to an unlucky random initialization of weights, we trained the model 100 times on one specific train-val-test split that had done well once in a previous iteration. Table 8 contains these results.

Table 8. Results based on 100 iterations with a fixed train-val-test split

Lead	v6, vz
Average (%)	99.89
Std. Dev. (%)	0.43
Median (%)	100

Through this exercise, we found a few train-val-test split cases in which the classification accuracy was suboptimal. The results on train-val-test splits with different distributions are poor because of the nature of the splits and are unable to assess the model’s capabilities, which is what we filter out. This could be remedied with access to a larger dataset to further improve the resilience and robustness of our neural network model; we hope to investigate this in the near future. A key goal of this project is to engage with the cardiology community to further improve the performance of the models introduced herein by getting access to larger data banks, and to design ready-to-use apps based on deep learning algorithms for their use in realistic diagnostics settings.

To put our results in context, we have compiled contemporaneous results in the literature, which have utilized the same dataset, along with ours in Table 9.

¹ A poor train-val-test split refers to one where the distribution of the train data does not match the distribution of the test data. With such a small dataset, this occurs at a relatively high frequency when the data is split up randomly.

Table 9. Comparison of our results with other studies in the literature that also use the PTB database. Best results are marked in boldface.

Work	Accuracy (%)	Sensitivity (%)	Specificity (%)	Precision (%)
Acharya et al.	93.5	93.7	—	92.8
Safdarian et al.	94.7	—	—	—
Kojuri et al.	95.6	93.3	—	97.9
Sun et al.	—	92.6	—	82.4
Liu et al.	94.4	—	—	—
Sharma et al.	96	93	—	99
Kachuee et al.	95.9	95.1	—	95.2
Remya et al.	93.61	93.22	94.28	—
Reasat et al.	84.54	85.33	84.09	—
Zewdie et al.	98.3	98.7	96.4	—
Feng et al.	95.4	98.2	86.5	—
Strodthoff et al.	—	93.3	89.7	—
Huang et al.	96.96	99.89	92.51	95.35
Liu et al.	98.59	99.53	94.50	—
Ours	99.43	99.40	99.45	99.46

4.2 Patient-Wise Split

The dataset used contained 549 records from 290 unique patients. This indicates that some patients had multiple records (i.e., multiple diagnoses). It is important to note that these diagnoses could have had different outcomes (a given patient could be healthy in one record and unhealthy in another). The analysis conducted above was done on a record-wise split, meaning that if a record of a particular patient was in the train set, a different record of that same patient could show up in the validation or test set. We investigated how our results would change if we instead did a patient-wise split; that is, all records of a particular patient show up only in either the train, validation, or test split. The procedure used for a record-wise split was replicated for a patient-wise split – the results for the best performing pairs of channels are presented in Tables 10 and 11.

Table 10. Quantification of accuracies for pairs of channels (patient-wise split)

Leads	i, vz	ii, v1	ii, vz	avr, v6	avf, vz	v1, v6
Accuracy (%)	93.93	92.94	95.73	94.51	93.01	91.76
Std. Dev. (%)	0.96	1.33	1.19	0.84	1.17	1.59
Median (%)	93.91	92.89	95.59	94.54	92.77	92.08

Table 11. As Table 10, continued

Leads	v4, vz	v5, v6	v5, vx	v6, vx	v6, vz	vy, vz
Accuracy (%)	93.18	94.15	94.45	94.28	96.47	94.33
Std. Dev. (%)	0.90	0.90	0.73	1.12	0.51	0.91
Median (%)	92.91	94.69	94.70	94.26	96.23	94.17

As can be seen from the Tables 10 and 11, the best performing pair of channels with a patient-wise split is (v6, vz), which happens to be the same as what we discovered for a record-wise split. The second best pair, (ii, vz) also contains the vz channel. We decided to see if combining these three channels (v6, vz, ii) would give us better results. After conducting a study of these three channels with various train-val-test splits, we found that this combination does yield better results. The following table presents our results over 10 trials (Table 12):

Table 12. Results for three channels: (v6, vz, ii)

Lead	v6, vz, ii
Average (%)	97.83
Std. Dev. (%)	0.56
Median (%)	97.92

As evident, using a combination of three channels is better than using pairs of channels. The patient-wise split results obtained here are lower than the results obtained for a record-wise split – this is expected as in a record-wise split, there may be some leakage from the train data into the test data. The other studies we compared our results to in Table 9 do not explicitly mention whether a record-wise or patient-wise split was used.

5 Discussion

This paper introduces a new architecture for heart condition classification based on raw ECG signals using multiple leads. Our approach outperforms the current state-of-the-art model on this dataset by a large margin of almost one full percent, with no manual feature engineering.

Another way in which our approach differs from others in the community is in our choice of leads. While other works in this domain have employed simply one or two of the 15 leads that the PTB database provides for each record, strong justification for their choice of lead(s) hasn't been provided.

Here, we studied which of the 15 ECG channels (12 conventional ECG leads and 3 Frank leads) contains the most meaningful information with respect to

myocardial infarction detection, finding that channels v6, vz and ii are the most significant. Our work also illustrates that recent advances in machine learning can be leveraged to produce a model capable of classifying myocardial infarction with a cardiologist-level success rate.

Limitations. As in many other deep learning studies in the literature, one of the key limitations to maximize the power of neural network models is the lack of labelled data [40]. We also face a similar challenge in this study, and candidly acknowledge that to further improve the performance of our models we require access to a larger, labelled data set. With a small database such as the PTB database, it is difficult to gauge the true robustness of the model since the test set ends up being relatively small. We are in the search for a larger data set to further increase the accuracy and generalization of our models, e.g., to correctly identify other heart conditions beyond myocardial infarction. If that were accomplished, then we would readily share our deep learning models and develop an open source, computationally efficient app that may be readily used by cardiologists.

6 Conclusion

We have conducted a detailed analysis of the relative importance of each of the standard 15 ECG channels to identify myocardial infarction using a multi-channel neural network model. Our findings indicate that deep learning can identify this heart condition upon processing only ten seconds of raw ECG data from the v6, vz and ii leads, reaching cardiologist-level success rate.

Through this analysis we: (i) furnished evidence that deep learning algorithms may be readily used as commodity software, i.e., taking two disparate areas of research—cardiology and seismology—we showed that a neural network model that was originally designed to identify earthquakes by processing multiple input channels may be re-designed and tuned to identify myocardial infarction; (ii) deep learning does not require feature engineering of ECG data to identify myocardial infarction in the PTB database. Indeed, our deep learning model only requires ten seconds of raw ECG data to identify this heart condition with cardiologist-level performance.

We also recommend to provide deep learning researchers access to a larger database to further improve and extend to other types of heart conditions the results presented in this manuscript. We look forward to working with the cardiology community to develop deep learning algorithms that may be readily applicable in a diagnostics setting, providing trustworthy, real-time information regarding heart conditions with minimal computational resources that are either deployed on the cloud or in lightweight, personalized devices.

Acknowledgments. EAH and ZZ gratefully acknowledge National Science Foundation (NSF) awards OAC-1931561 and OAC-1934757. AG acknowledges support from

the Fiddler Innovation Undergraduate Fellowship and the Students Pushing Innovation (SPIN) program at the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign.

This work utilized resources supported by the NSF’s Major Research Instrumentation program, grant OAC-1725729, as well as the University of Illinois at Urbana-Champaign. We are grateful to NVIDIA for donating several Tesla P100 and V100 GPUs that we used for our analysis, and the NSF grants NSF-1550514, NSF-1659702 and TG-PHY160053. This research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. We thank the NCSA Gravity Group for useful feedback.

References

1. Acharya, U.R., et al.: Application of deep convolutional neural network for automated detection of myocardial infarction using ECG signals. *Inf. Sci.* **415**, 190–198 (2017). <https://doi.org/10.1016/j.ins.2017.06.027>
2. Ackermann, S.: Using transfer learning to detect galaxy mergers. *Mon. Notices Roy. Astron. Soc.* **479**(1), 415–425 (2018). <https://doi.org/10.1093/mnras/sty1398>
3. Gupta, A.: Myocardial Infarction Detection. <https://github.com/arjung128/mi-detection>
4. Berg, J., Nyström, K.: A unified deep artificial neural network approach to partial differential equations in complex geometries. *Neurocomputing* **317**, 28–41 (2018)
5. Bhattacharya, T., et al.: AI meets exascale computing: Advancing cancer research with large-scale high performance computing. *Front. Oncol.* **9**, 984 (2019). ISSN 2234-943X. <https://doi.org/10.3389/fonc.2019.00984>, <https://www.frontiersin.org/article/10.3389/fonc.2019.00984>
6. Choma, N., Monti, F., et al.: Graph neural networks for icecube signal classification. In: *Proceedings - 17th IEEE International Conference on Machine Learning and Applications, ICMLA 2018*, vol. 1, pp. 386–391. Institute of Electrical and Electronics Engineers Inc. (2019). <https://doi.org/10.1109/ICMLA.2018.00064>
7. Chua, A.J., Galley, C.R., Vallisneri, M.: Reduced-order modeling with artificial neurons for gravitational-wave inference. *Phys. Rev. Lett.* **122**(21), 211101 (2019). <https://doi.org/10.1103/PhysRevLett.122.211101>
8. Dominguez Sánchez, H., et al.: Improving galaxy morphologies for SDSS with deep learning. *Mon. Notices Roy. Astron. Soc.* **476**, 3661–3676 (2018). <https://doi.org/10.1093/mnras/sty338>
9. Dominguez Sánchez, H., et al.: Transfer learning for galaxy morphology from one survey to another. *Mon. Notices Roy. Astron. Soc.* **484**(1), 93–100 (2019). <https://doi.org/10.1093/mnras/sty3497>
10. Feng, K., et al.: Myocardial infarction classification based on convolutional neural network and recurrent neural network. *Appl. Sci.* (2019). <https://doi.org/10.3390/app9091879>
11. George, D., Huerta, E.A.: Deep Learning for real-time gravitational wave detection and parameter estimation: Results with AdvancedLIGO data. *Phys. Lett. B* **778**, 64–70 (2018). <https://doi.org/10.1016/j.physletb.2017.12.053>. [arXiv: 1711.03121](https://arxiv.org/abs/1711.03121) [gr-qc]
12. George, D., Huerta, E.A.: Deep neural networks to enable real-time multimessenger astrophysics. *Phys. Rev. D* **778**, 97–044039 (2018). <https://doi.org/10.1103/PhysRevD.97.044039>. [arXiv: 1701.00008](https://arxiv.org/abs/1701.00008) [astro-ph.IM]

13. George, D., Shen, H., Huerta, E.A.: Classification and unsupervised clustering of LIGO data with deep transfer learning. *Phys. Rev. D* **97**, 101501 (2018). <https://doi.org/10.1103/PhysRevD.97.101501>. [arXiv: 1706.07446](https://arxiv.org/abs/1706.07446) [gr-qc]
14. Goldberger, A.L., et al.: PhysioBank, physiotookit, and physionet: components of a new research resource for complex physiologic signals. *Circulation* **101**(23), e215–e220 (2000)
15. Huerta, E.A., et al.: Convergence of artificial intelligence and high performance computing on NSF-supported cyberinfrastructure, March 2020. *arXiv e-prints*, [arXiv:2003.08394](https://arxiv.org/abs/2003.08394) [physics.comp-ph]
16. Huerta, E.A., et al.: Enabling real-time multi-messenger astrophysics discoveries with deep learning. *Nat. Rev. Phys.* **1**(10), 600–608 (2019). <https://doi.org/10.1038/s42254-019-0097-4>. [arXiv:1911.11779](https://arxiv.org/abs/1911.11779) [gr-qc]
17. Fawaz, H.I., et al.: Deep learning for time series classification: a review. *Data Mining Knowl. Disc.* **33**(4), 917–963 (2019). <https://doi.org/10.1007/s10618-019-00619-1>. ISSN 1573-756X
18. Kachuee, M., Fazeli, S., Sarrafzadeh, M.: ECG heartbeat classification: a deep transferable representation. In: *IEEE Conference on Healthcare Informatics*, July 2018. [arXiv: 1805.00794](https://arxiv.org/abs/1805.00794)
19. Asad Khan, E.A., et al.: Deep learning at scale for the construction of galaxy catalogs in the dark energy survey. *Phys. Lett. B* **795**, 248 – 258 (2019). ISSN 0370-2693. <https://doi.org/10.1016/j.physletb.2019.06.009>. <http://www.sciencedirect.com/science/article/pii/S0370269319303879>
20. Kojuri, J., et al.: Prediction of acute myocardial infarction with artificial neural networks in patients with non-diagnostic electrocardiogram. *J. Cardiovasc. Dis. Res.* (2015). <https://doi.org/10.5530/jcdr.2015.2.2>
21. Lecun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**, 436–444 (2015). <https://doi.org/10.1038/nature14539>
22. Liu, B., et al.: A novel electrocardiogram parameterization algorithm and its application in myocardial infarction detection. *Comput. Biol. Med.* (2015). <https://doi.org/10.1016/j.combiomed.2014.08.010>
23. Liu, N., et al.: A simple and effective method for detecting myocardial infarction based on deep convolutional neural network. *J. Med. Imaging Health Inform.* (2018). <https://doi.org/10.1166/jmih.2018.2463>
24. T. Perol, M. Gharbi, and M. Denolle. Convolutional neural network for earthquake detection and location. *Sci. Adv.* **4**, e1700578 (2018). <https://doi.org/10.1126/sciadv.1700578>, pmid: 29487899
25. Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686 – 707 (2019). ISSN 0021-9991. <https://doi.org/10.1016/j.jcp.2018.10.045>. <http://www.sciencedirect.com/science/article/pii/S0021999118307125>
26. Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics informed deep learning (part I): data-driven solutions of nonlinear partial differential equations. *CoRR*, abs/1711.10561 (2017). <http://arxiv.org/abs/1711.10561>
27. Reasat, T., Shahnaz, C.: Detection of inferior myocardial infarction using shallow convolutional neural networks. In: *IEEE Region 10 Humanitarian Technology Conference*, August 2018. <https://doi.org/10.1109/R10-HTC.2017.8289058>
28. Rebei, A., et al.: Fusing numerical relativity and deep learning to detect higher-order multipole waveforms from eccentric binary black hole mergers. *Phys. Rev. D* **100**(4), 044025 (2019). <https://doi.org/10.1103/PhysRevD.100.044025>

29. Remya, R.S., Indiradevi, K.P., Anish Babu, K.K.: Classification of myocardial infarction using multi resolution wavelet analysis of ECG. In: International Conference on Emerging Trends in Engr., Science and Technology, December 2016. <https://doi.org/10.1016/j.protcy.2016.05.195>
30. Safdarian, N., Dabanloo, N.J., Attarodi, G.: A new pattern recognition method for detection and localization of myocardial infarction using t-wave integral and total integral as extracted features from one cycle of ECG signal. J. Biomed. Sci. Eng. (2014). <https://doi.org/10.4236/jbise.2014.710081>
31. Schmidhuber, J.: Deep learning in neural networks: an overview. Neural Netw. **61**, 85 – 117 (2015). ISSN 0893-6080. <https://doi.org/10.1016/j.neunet.2014.09.003>. <http://www.sciencedirect.com/science/article/pii/S0893608014002135>
32. Sharma, L.D., Sunkaria, R.K.: Inferior myocardial infarction detection using stationary wavelet transform and machine learning approach. Signal Image Video Process. (2017). <https://doi.org/10.1007/s11760-017-1146-z>
33. Shen, H., et al.: Denoising gravitational waves with enhanced deep recurrent denoising auto-encoders, pp. 3237–3241, May 2019. <https://doi.org/10.1109/ICASSP.2019.8683061>
34. Shen, H., Huerta, E.A., Zhao, Z.: Deterministic and Bayesian Neural Networks for Low-latency Gravitational Wave Parameter Estimation of Binary Black Hole Mergers (2019). [arXiv: 1903.01998](https://arxiv.org/abs/1903.01998) [gr-qc]
35. Springenberg, J.T., et al.: Bayesian optimization with robust bayesian neural networks, pp. 4134–4142. Curran Associates, Inc. (2016)
36. Sun, L., et al.: ECG analysis using multiple instance learning for myocardial infarction detection. IEEE Trans. Biomed. Eng. (2012). <https://doi.org/10.1109/TBME.2012.2213597>
37. Tang, H., Scaife, A.M.M., Leahy, J.P.: Transfer learning for radio galaxy classification. Mon. Notices Roy. Astron. Soc. **488**(3), 3358–3375 (2019). <https://doi.org/10.1093/mnras/stz1883>
38. United States Department of Energy. Scientific Machine Learning. <https://www.osti.gov/servlets/purl/1478744>
39. Wei, W., Huerta, E.A.: Gravitational wave denoising of binary black hole mergers with deep learning. Phys. Lett. B **800**, 135081 (2020). ISSN 0370-2693. <https://doi.org/10.1016/j.physletb.2019.135081>. <http://www.sciencedirect.com/science/article/pii/S0370269319308032>
40. Wei, W., et al.: Deep transfer learning for star cluster classification: I. application to the PHANGS-HST survey. Mon. Notices Roy. Astron. Soc. **493**(3), 3178–3193 (2020). <https://doi.org/10.1093/mnras/staa325>
41. Zewdie, G., Xiong, M.: Fully automated myocardial infarction classification using ordinary differential equations. *arXiv e-prints*, [arXiv:1410.6984](https://arxiv.org/abs/1410.6984), October 2014