

Impact of Neural Architecture Design on Cardiac Abnormality Classification Using 12-lead ECG Signals

Najmeh Fayyazifar¹, Selam Ahderom¹, David Suter¹, Andrew Maiorana^{2,3}, Girish Dwivedi^{3,4}

¹Edith Cowan University, Perth, Australia

²Curtin University, Perth, Australia

³Fiona Stanley Hospital, Perth, Australia

⁴University of Western Australia, Perth, Australia

Abstract

Cardiac rhythm abnormality, as associated with irregular heart activity, presents as changes in an electrocardiogram (ECG). In this paper, as part of the PhysioNet Challenge 2020, we propose two cardiac abnormality detection and classification neural models, using 12-lead ECG signals. Our ECU team proposes a hand-designed Recurrent Convolutional Neural Network (RCNN), consisting of 49 one-dimensional convolutional layers, 16 skip connections, and one Bi-Directional LSTM layer. This model, without relying on any pre-processing or manual feature engineering, achieved a Challenge validation score of 62.3% and a full test score of 38.2%, ranking us 9th out of 41 teams in the official ranking. Our second neural model, designed through neural architecture search, did not score on the full test dataset nor on the validation dataset; however, we optimistically expect performance improvement compared to our hand-designed RCNN model. This model scored 64.4% using 10-fold cross-validation on the training dataset, which is 2.5% higher than the training score of our RCNN model, using 10-fold cross-validation.

1. Introduction

Cardiovascular disease was reported to be the leading cause of death worldwide in 2012 (accounting for 17.3 million deaths per year), and has been projected to grow to 23.6 million deaths by 2030 [1]. According to the American Heart Association, cardiovascular disease causes 25% of annual deaths in the United States [2]. Development of an automatic cardiac abnormal activity detection system could assist medical personnel in providing timely and accurate diagnosis of cardiovascular disease related to rhythm abnormalities. The PhysioNet/Computing in Cardiology (CinC) Challenge 2020, provides an opportunity for the machine learning community to propose scientific solutions to the problem of automatic detection of cardiac abnormalities from standard 12-lead ECGs [3].

In the literature, most successful machine learning

based cardiac abnormality detection and classification methods have utilized manually engineered feature extraction schemes [4]. Whilst, these techniques when combined with classical classifiers provide promising classification performance, they require expert knowledge to manually design feature extraction methods. In this paper, we study the role of neural architecture design on the classification of cardiac abnormalities from ECG signals. More specifically, we propose two neural models that produce promising results, without requiring feature extraction and/or feature engineering. The first model is constructed by combination of a hand-designed Convolutional Neural Network (CNN) and a Long-Short Term Memory (LSTM) model. Although this model has provided competitive results on the full Challenge test dataset (ranked 9th), the manual design of the network's structure is the result of a time-demanding task that requires expert knowledge in deep learning optimization. In our second proposed model, we automate the process of neural architecture design and construct the architecture of this model with minimal reliance on human expert knowledge.

The rest of this paper is organized as follows. In section 2 we present our proposed neural architectures while in section 3 we demonstrate our experimental results. Finally, in section 4 we discuss our results and conclusions.

2. Method

2.1. Dataset and data preparation

The dataset used in this study, as provided by the PhysioNet Challenge 2020, has 43101 multi-label ECG samples containing 111 classes of heartbeat rhythms [3]. The performance metric defined by Challenge organizers uses a subset of 27 of these classes for the scoring of competing algorithms, thus samples that belonged to unused classes were discarded from our training dataset, reducing the total number of training samples to 37619. Amongst the 27 classes, 3 pairs of classes were considered equivalent for scoring performance, hence we treated each pair as the same class, leading the final classification problem to be a 24-class, multi-label classification task.

The length of ECG samples varies between 6 seconds to 30 minutes, with a median of 10 seconds. As CNN models accept signals with a constant length, waveforms shorter than 10 seconds were zero-padded, and signals longer than 10 seconds were truncated to the last 10 seconds.

In order to evaluate the performance of our proposed models, we have used 10-fold cross validation on the training dataset, that is, we have randomly partitioned the training dataset into 10 equally sized subsets, then trained our models 10 times using one subset for testing, 8 for training as well as one for validation (for determination of early stopping), and reported mean of 10 held-out folds of the training data.

In order to reduce the size of input to our NAS-based model without losing meaningful information, and consequently reducing the amount of required computing resources such as memory, the input time domain signal was transformed into wavelet presentation. This was completed by using Db1 mother wavelet with one level of decomposing, providing a set of details and approximation coefficients. In the corresponding experiments, approximation coefficients were used as the input to our model while details were discarded. Our main intention for using wavelet transformation was to reduce the required computing resources; however, Christov et al. [5] showed that wavelet transforms are one of the most suitable methods for ECG analysis. Accordingly, we expect performance improvement.

2.2. Our proposed RCNN architecture

Recurrent Neural Networks (RNNs) which model the temporal dependencies of data, have proved to perform well for time series classification [6]. The heart’s electrical conduction system follows a periodic process; thus, such temporal dependencies naturally exist in ECG signals [7]. In this paper, we have experimentally shown that our proposed RCNN model, constructed by combination of a proper CNN model, that effectively extract features from ECG waveforms, and a Long-Short Term Memory (LSTM) model, that models the dependencies in ECG waveforms, can provide comparable results to methods utilizing extensive manual feature engineering. Our proposed architecture is illustrated in Figure 1. Inspired by architecture proposed in [8], the CNN section of our proposed architecture is constructed of 16 blocks with three 1-Dimensional convolution layers within each block. All convolutions are followed by a Batch Normalization (BN) layer, a Rectifier Linear Unit (ReLU) activation function, and a dropout layer. All blocks have been connected to their previous block by a skip connection, similar to that of Residual Networks (ResNet) [9]. These skip connections concatenate features at the end of the last two blocks, allowing information to propagate through the deep CNN structure, thus reducing the effect of vanishing gradient descent. Prior to our first block, the 12-lead ECG




Figure 1. Our proposed RCNN model

inputs were passed through a stem convolution layer with 32 filters, expanding the total number of convolution layers in our model to 49. The number of output channels has been doubled at the end of the 4th, 8th, and 12th blocks, leading to a total of 256 output channels at the last convolution layer. The length of all convolution filters is set to 24. Within each block, the first 2 convolution layers have a stride of 1, and the last layer’s stride is 2, leading the size of input signal to be divided by 2 by the end of each block. At each skip connection, a max-pooling layer with a stride of 2 was applied to make the output size of the previous block consistent with output size of current block (to allow concatenation). The output of the last convolution layer is then passed through a BiLSTM layer. The number of hidden units of the BiLSTM layer is empirically adjusted to 300 where the evaluation metric was most promising. The last layer of the network is a sigmoid layer with 24 nodes (corresponds to 24 output classes) which predicts a value between 0 and 1 for each class in our multi-label classification problem, indicating if a sample belongs to a class. The threshold in a binary classification problem is usually 0.5; however, as studied in [10], if the classification problem is imbalanced this threshold can vary. In this paper, by applying moving decision threshold [11], we selected 0.2 as the classification threshold for testing our proposed model.

2.3. Our NAS-based architecture

In recent years, Neural Architecture Search (NAS) methods have been proposed to assist the machine learning community to automatically design suitable architectures yielding higher performance for a given problem. However, to date, there are very few studies that have applied a NAS method to design a CNN model for biomedical signal classification. In a previous study [12], we explored the performance of a well-known NAS method, Efficient Neural Architecture Search (ENAS) [13], on Atrial Fibrillation (AF) detection, confirming that NAS-Based methods can provide competitive results compared to state of the art methods for biomedical signal classification tasks.

In this paper, we proposed an automatically designed CNN architecture by employing Differentiable Architecture Search (DARTS) method [14], modifying its search space to tailor it to 1-Dimensional ECG waveforms.

This approach consists of two phases. In first phase, the search algorithm looks for the best possible set of operations within the search space and builds up the best possible architecture. In the second phase, the discovered architecture is trained from scratch to provide the final neural model. Our search space consists of: a) 1-dimensional convolutions (with possible filter lengths: 3,5,7,9,11), b) dilated convolution (filter length 3,5), c) max-pooling (size of 3), d) average-pooling (size of 3), and e) skip connections. Convolution operations have a ReLU-Conv-BN order. The DARTS search method considers each of these blocks as building units of a CNN architecture and optimizes possible operations within the block, whilst forming the final architecture by stacking these blocks.

In the search algorithm, the total number of blocks were set to 8 where block numbers 3 and 6 are reduction blocks. In reduction blocks the number of convolution channels are doubled while input signal length is reduced by a factor of two. Each of our proposed blocks has two input nodes (outputs from two previous blocks or input data in the case of first two blocks), two middle nodes, and one output node that concatenates the outputs of the middle nodes. Each middle node can be connected to any two previous nodes. Similar to our RCNN model, the last layer is a sigmoid layer with 24 nodes.

The search algorithm in DARTS uses a bi-level optimization strategy, as formulated in Equation 1. In this equation, W , the weights of the network, are optimized by minimizing training loss, while α , the selected architectures, are optimized by minimizing validation loss. Each epoch of the search algorithm samples a neural architecture and evaluates its performance by computing its validation loss. This algorithm moves towards an optimized architecture by minimizing validation loss. We have chosen the architecture with minimum validation loss as the final architecture and trained it from scratch. The discovered normal and reduction blocks are illustrated in Figures 2.

2.4. Experimental setup

The parameters of the Our RCNN model are initialized with the ‘He’ initializer. An Adam optimizer was used to train the model parameters. The learning rate was initialized at 10^{-3} and a “ReduceLROnPlateau” learning scheduler with a decay rate of 0.1 and a patience of 3 was used. The minimum learning rate was set to 10^{-6} . The algorithm was trained for 50 epochs, and early stopping with patience of 15 was set. On average, each training epoch of the algorithm took about 300s, leading to a maximum training time of 250 minutes on a single Nvidia Geforce GTX 1080 Ti GPU.

The parameters of our NAS-based network were trained using an SGD optimizer with momentum of 0.9. The learning rate was initiated at 0.025, and a “cosine




Figure 2. a) Our discovered normal block. b) our discovered reduction block.

annealing” learning rate scheduler with a minimum learning rate of 10^{-3} was used. The search algorithm was run for 50 epochs and the architecture with minimum validation loss was chosen for training from scratch. Each epoch of search algorithm took about 50 minutes on a single Nvidia Quadro RTX 8000 GPU. The training from scratch was run for 200 epochs using early stopping with a patience of 20. The train algorithm run time was reported 20 hours on the same Quadro GPU machine.

$$\min_{\alpha} L_{val}(W^*(\alpha), \alpha)$$

$$s.t. W^*(\alpha) = \arg \min_W L_{train}(W, \alpha)$$

Equation 1.

3.1. Results - RCNN model

We trained our proposed Conv-BiLSTM model using time domain ECG signal. The derived model was evaluated using 10-fold cross validation on training dataset, where it received a score of 61.9%. This model was submitted to the Challenge and received the validation score of 62.3% and full test score of 38.2%, placing our ECU team 9th in the official ranking. To study the effect of the LSTM model and compare it with the base CNN model, we trained and evaluated the CNN model with the same setting, except removing the BiLSTM section and decreasing our initial learning rate to 5×10^{-4} for the CNN model as with our default learning rate the network did not converge. The evaluation score on training dataset using 10-fold cross validation was 59.2% which shows the importance of BiLSTM model.

In order to study the effect of wavelet transformation, we have computed the wavelet coefficients, fed them as input to our Conv-BiLSTM model, and evaluated the results. Our proposed network performance is 62.9% using 10-fold cross validation on training dataset. Whilst this model has not received a score on validation dataset, we estimate some improvements over our scored model.

3.2. Results - NAS-based model

In these experiments, wavelet transformations of ECG waveforms were fed as input to the search algorithm. We performed the search algorithm on our designed search space. We chose the normal and reduction blocks (shown in Figure 2.) sampled at the last epoch of the algorithm where the validation loss is minimum and trained the architecture constructed by stack of those blocks (20 blocks where block numbers 7 and 14 are reduction). The

performance of our model was reported as 64.4% using 10-fold cross validation on training dataset. Whilst this model failed to receive a score on the validation dataset and the full test dataset, we expect a roughly 2-3% improvement compared to our hand-designed model, which might have led us to a higher place in the Challenge’s official ranking. Table 1 summarizes all our proposed models together with their score on training dataset using 10-fold cross validation, Challenge’s validation dataset, and full test dataset.

Table 1. Summary of our proposed models’ score

Model	Wavelet decomp.	10-fold cross validation (training dataset)	Validation set score	Full test dataset
Hand-designed CNN	No	59.2%	-----	-----
Hand-designed Conv-BiLSTM	No	61.9%	62.3%	38.2%
Hand-designed Conv-BiLSTM	Yes	62.9%	-----	-----
NAS-based CNN	Yes	64.4%	-----	-----

4. Discussion and conclusion

Early detection of cardiac rhythm abnormalities can improve the quality of treatments that patients receive. In this study, we have proposed two neural models which can be used for early diagnosis of 24 classes of cardiac abnormalities without relying on manual feature engineering methods. Our experiments show that proper neural architecture design has a significant impact on model classification performance. By incorporating long-term and short-term memory, BiLSTM networks can learn temporal features and dependencies more accurately than a hand-designed CNN, leading to higher performance. The forget gate in the LSTM models allow the model to discard parts of ECG waveforms that don’t contribute towards classification. Long-term memory facilitates the learning process by remembering discriminative features from all parts of the waveform. Our Conv-BiLSTM model reported a score of 61.9% on training dataset using 10-fold cross validation, which is 2.7% higher than our hand-designed CNN model on the same set. The Conv-BiLSTM model received 62.3% score on validation dataset and 38.2% on full test dataset, placed our team 9th in official ranking.

Pre-processing techniques such as wavelet decomposition can boost classification performance. We would expect it to marginally improve our score on the validation and full test datasets. With this technique, our experiments suggest a 1% higher classification score on training dataset using 10-fold cross validation, reaching 62.9%. Our experiments indicate that the most significant performance boost can be derived by using a NAS method to search over a subset of standard operations in CNNs, and automatically design the best possible CNN model. Our NAS-based neural model reported a score of 64.4% using

10-fold cross validation on training dataset. This model did not receive a score on validation and full test datasets; however, we expected around a 3% improvement compared to our current score in official ranking. We acknowledge that this conclusion might be optimistic.

In future work, it will be desirable to combine LSTM networks and automatically designed CNN architectures, which might provide more promising classification results. Moreover, the search space of our NAS algorithm can be further modified, where the effects of different search space designs on model performance can be further explored.

References

- [1] Laslett, L.J., Alagona, P., Clark, B.A., et al.: ‘The worldwide environment of cardiovascular disease: prevalence, diagnosis, therapy, and policy issues: a report from the American College of Cardiology’, *Journal of the American College of Cardiology*, 2012, 60, (25 Supplement), pp. S1-S49
 - [2] Benjamin, E.J., Muntner, P., Alonso, A., et al.: ‘Heart disease and stroke Statistics-2019 update a report from the American Heart Association’, *Circulation*, 2019
 - [3] Alday, E.A.P., Gu, A., Shah, A., Robichaux, C., et al.: ‘Classification of 12-lead ECGs: the PhysioNet/Computing in Cardiology Challenge 2020’, *Physiological Measurement*, 2020
 - [4] Teijeiro, T., Félix, P., Presedo, J., and Castro, D.: ‘Heartbeat classification using abstract features from the abductive interpretation of the ECG’, *IEEE Journal of Biomedical and Health Informatics*, 2016, 22, (2), pp. 409-420
 - [5] Christov, I., Gómez-Herrero, G., Krasteva, V., et al.: ‘Comparative study of morphological and time-frequency ECG descriptors for heartbeat classification’, *Medical Engineering & Physics*, 2006, 28, (9), pp. 876-887
 - [6] Hüsken, M., and Stagge, P.: ‘Recurrent neural networks for time series classification’, *Neurocomputing*, 2003, 50, pp.223-235
 - [7] Seymour, R.S., Hargens, A.R., and Pedley, T.J.: ‘The heart works against gravity’, *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 1993, 265, (4), pp. R715-R720
 - [8] Hannun, A.Y., Rajpurkar, P., Haghpanahi, M., et al.: ‘Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network’, *Nature Medicine*, 2019, 25, (1), pp. 65
 - [9] He, K., Zhang, X., Ren, S., et al.: ‘Deep residual learning for image recognition’, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2016.), pp. 770-778
 - [10] He, H., and Ma, Y.: ‘Imbalanced learning: foundations, algorithms, and applications’ (*John Wiley & Sons*, 2013.)
 - [11] Provost, F.: ‘Machine learning from imbalanced data sets 101’, (*AAAI Press*, 2000), pp. 1-3
 - [12] Fayyazifar, N.: ‘An accurate CNN Architecture for Atrial Fibrillation Detection Using Neural Architecture Search’, *Proceedings of 28th European Signal Processing Conference* (2020), pp.1135-1139
 - [13] Pham, H., Guan, M.Y., Zoph, B., et al.: ‘Efficient neural architecture search via parameter sharing’, *arXiv preprint arXiv:1802.03268*, 2018
 - [14] Liu, H., Simonyan, K., and Yang, Y.: ‘Darts: Differentiable architecture search’, *arXiv preprint arXiv:1806.09055*, 2018
- Address for correspondence: Najmeh Fayyazifar, Fayyazifar1@ecu.edu.au