

Deep Discriminative Domain Generalization with Adversarial Feature Learning for Classifying ECG Signals

Zuogang Shang¹, Zhibin Zhao¹, Hui Fang², Samuel Relton³, Darcy Murphy⁴, Zoe Hancox³, Ruqiang Yan¹, David Wong⁴

¹Xi'an Jiaotong University, Xi'an, China

²Loughborough University, Loughborough, UK

³University of Leeds, Leeds, UK

⁴University of Manchester, Manchester, UK

Abstract

Introduction: The goal of the 2021 PhysioNet/CinC challenge is to classify cardiac abnormalities from ECGs and evaluate the diagnostic potential of reduced-lead ECGs. Here, we describe the classification model created by the team "AI_Healthcare".

Methods: ECGs were downsampled to 300 Hz and filtered by wavelet. ECGs were randomly clipped or zero-padded to 4,096 samples. We modified an SE-ResNet to perform multi-task classification of both dataset and disease. We used a gradient reversal layer as part of an adversarial feature learning scheme to learn domain-invariant and discriminative representations.

Results: We trained our domain-invariant model on 5 datasets, keeping one data set (Ningbo) for local validation. We also trained a baseline SE-ResNet using the same training data. In validation on the held-out data set, the domain-invariant model had a higher Challenge metric than the baseline model. Our entry was not officially ranked in the Challenge, as we did not have a successful entry during the unofficial phase of the Challenge.

Conclusion: The domain-invariant model performed better than the baseline model in local held-out datasets, suggesting that this method may help improve generalisation performance.

1. Introduction

In this paper we describe a deep learning model developed to classify cardiac abnormality from 12-lead, 6-lead, 4-lead, 3-lead and 2-lead electrocardiogram (ECG) signals with varying sample lengths and frequencies. 12-lead ECGs are used clinically to diagnose cardiac abnormalities by measuring the electrical activity of the heart. Reduced-lead ECGs are also being explored for their diagnostic potential to reduce recording time and expense, and improve

ease of use in clinical settings [1].

ECG classification using deep learning models, such as the one described in this paper, may have the ability to automatically diagnose a range of cardiac abnormalities without requiring all 12-leads, which could reduce resource demand. However, in the Physionet 2020 challenge [2], all models suffered from poorer performance on a hidden dataset from an undisclosed location. We aim to address this issue by building on the work of the previous deep neural network architecture [3], incorporating domain generalisation through adversarial feature learning.

2. Methods

Our goal was to create an ECG classification model that learned domain-agnostic features, and that could also be applied to reduced-lead ECGs. We used a modified ResNet with a Squeeze-and-Excitation (SE) attention block [4] to extract deep features. Combined with hand-crafted features, a multi-source adversarial network was trained to learn useful domain-invariant features for the main task of diagnosing cardiac abnormalities. We expected that the domain-invariant representation would perform worse on the test data from the seen datasets. As the seen datasets have a frequency of 500Hz and the majority of the unseen test data is of a different frequency, we used a model without domain generalisation (baseline model) for test examples with a frequency of 500Hz.

2.1. Data Pre-processing

The public datasets were CPSC [5], CPSC-Extra, INCART [6], PTB [7], PTB-XL [8], G12EC, Chapman-Shaoxing [9], and Ningbo [10]. Four pairs of the 30 scored classes were considered equivalent, making this effectively a 26-class problem.

ECGs were resampled to 300 Hz for input to the deep

model. During the training phase, we chose a signal length of 4096 samples. Shorter signals were randomly zero-padded and longer signals were randomly clipped.

To reduce unwanted noise, we employed wavelet denoising [11]. ECGs were decomposed into 9 levels with Daubechies D_6 ('db6') wavelet. We replaced the first approximation sub-band (baseline wander) and the first detail sub-band (little relevant information) with zeros. The other detail sub-bands were used to reconstruct the signal.

Age, gender, and Heart Rate Variability (HRV) features were concatenated with deep features. Unknown values of age and gender were masked and set to 0.

HRV features were extracted from lead I and II. First, R peak locations were extracted using the EngZee QRS detector [12]. These peaks were used to derive: the standard deviation of R-peak, or normal-to-normal, (SDNN) intervals, root mean square of successive R-peak differences (RMSSD), the standard deviation of the successive differences (SDSD) between adjacent R-peak (NN) intervals, the proportion of NNs that are greater than 20 ms (NN20) divided by total number of R-peak intervals (PNN20), and heart-rate (HR). For normalization, SDSD was divided by 1000 and HR was divided by 100. All HRV features were set to zero if fewer than 5 R-peaks were detected in an example. Age, gender, and HRV features were encoded to a total of 17 feature values.

2.2. Model Description

Our model was designed to extract discriminative domain-invariant features from the input signals and extra features. It then uses the features to classify the ECG recordings into 26 classes. It achieves this by multi-task learning of ECG abnormalities and domain, with a loss function that seeks to maximise domain loss, and minimise ECG abnormality loss. The model is illustrated in Fig 1.

2.2.1. ResNet Feature Extraction

The initial branch of the model is a modified ResNet model with an adaptive input channel. The modified ResNet model from [3] consists of one convolution layer with a wide kernel and 8 residual blocks (RBs).

A wide kernel has been shown to perform well for time sequence classification [13]. We employ a convolution kernel size of 15 in the first layer followed by batch normalisation (BN) and a rectified linear unit (ReLU). 64 kernels are used in the first convolution layer.

The RB consists of two convolution layers with BN and ReLU in between. A dropout layer with dropout rate of 0.2 is also inserted to alleviate overfitting. After the second convolution, a BN layer and a SE block are used, followed by a residual connection from RB input and a ReLU layer. A convolution kernel with size of 7 is employed in the RB.

The number of kernels for the RB are 64, 64, 128, 128, 256, 256, 512, and 512. The feature dimension is halved after the third, fifth, and seventh RB. The SE block acts to adaptively recalibrate channel-wise feature response and calculates channel importance by explicitly modelling the dependencies between channels. The SE block contains a global average pooling layer, a bottleneck with two fully connected (FC) layers around a ReLU layer, and a sigmoid layer. The reduction between the two FC layers is 16. 8 RBs are used to enlarge model receptive field and improve feature extraction ability. The residual connection confirms the training process stability [14].

After deep feature extraction, we concatenate the deep feature set with the encoded HRV, age, and gender features to a total dimension of 546.

The features are used for two tasks, domain classification and ECG abnormality classification.

2.2.2. Domain Classifier

Data from different domains (datasets) may have a shift in distributions and representations [15]. We envisage that the final classification decisions should be based on representations that are both discriminative for the main task (ECG abnormality classification) and invariant to the domain changes.

A discriminative domain-invariant representation requires mapping a domain-variant representation into a similar representation in different domains. We divided our training datasets into seven domains by their recording file name and gave each ECG recording a domain label. The domain classifier consists of a simple three-layer bottleneck FC classifier and a Gradient Reversal Layer (GRL). We label the loss for this branch as L_2 .

By minimizing the domain label prediction loss L_2 , the domain classifier is optimized to learn domain features from input features. The GRL means that the gradient for L_2 is reversed for the feature extraction part of the network, meaning that the feature extractor tries to maximise L_2 . This leads to the feature extractor learning features which give the least domain information.

2.2.3. Discriminative Classifier

Multi-label ECG abnormality classifications are created from the 546-dimension features by using two FC layers with a middle dimension of 256. The loss for the discriminative classifier is L_1 .

2.2.4. Training Setup

The training error for multi-label classification was average binary cross entropy (BCE) loss L_1 . For the adversarial domain classification task, the loss was cross entropy




Figure 1. Architecture of the proposed model.

L_2 . The final loss L is:

$$L = L_1 + \lambda L_2. \quad (1)$$

The weight parameter, λ , was set empirically at 0.05. For training, we chose 0.0003 as the initial learning rate with the Adam optimiser. It was reduced tenfold in the 20th epoch. The model was trained for a total of 30 epochs with batch size of 64.

The baseline model was trained with the L_1 loss only.

2.3. Model Evaluation

Thresholds for different classes should be different because of class imbalance. After training, we used the validation signals to search for the best thresholds for the models: (1) Thresholds were initialised to be same for all classes and then searched in the range [0,1] with a step 0.1 to get an approximate threshold; (2) Adjust approximate threshold for each class by searching with in steps of 0.01 when all other thresholds are fixed. Validation signals shorter than 4,096 were zero-padded. Longer signals were segmented into multiple patches with overlap of 256.

We expected that in testing most examples with a frequency of 500Hz would be from the same domain as the training data (CPSC, CPSC-Extra, and G12EC). We planned to use the baseline model for 500Hz examples and the domain-invariant model for other examples.

3. Local Testing

For local comparison, we trained the model on N-1 of the N training data sets, reserving the Ningbo dataset as a

| Lead | 12 | 6 | 4 | 3 | 2 |
|-------------------|------|------|------|------|------|
| B (seen domain) | 0.75 | 0.71 | 0.73 | 0.73 | 0.71 |
| D (seen domain) | 0.72 | 0.68 | 0.69 | 0.69 | 0.68 |
| B (unseen domain) | 0.43 | 0.46 | 0.46 | 0.44 | 0.45 |
| D (unseen domain) | 0.44 | 0.49 | 0.48 | 0.48 | 0.49 |

Table 1. Challenge metric scores for the baseline and domain invariant models for data in the seen (average 5-fold cross-validation of training data from CPSC, CPSC-Extra, PTB, PTB-XL, and G12EC) and unseen (Ningbo dataset) domains. B: baseline model. D: domain invariant model.

local test. 5-fold cross-validation results (seen domain), and results for the Ningbo dataset (unseen domain) are shown in Table 1. The domain generation model obtained a better performance in the unseen Ningbo dataset compared to the baseline model.

4. Results

Due to a technical error, the domain-invariant model, which is trained on all the training data, was not tested on the hidden test set. Detailed scores for local 5-fold cross validation and scores on hidden validation set are shown in Table 2.

Scores for baseline model, which is only trained on CPSC, CPSC-Extra and G12EC, are shown in Table 3.

5. Discussion and Conclusions

The model described here was designed to be generalisable to unseen data sets, thanks to its ability to ex-

| Leads | Training | Validation | Test | Ranking |
|-------|-------------------|------------|------|---------|
| 12 | 0.721 ± 0.001 | 0.637 | N/A | none |
| 6 | 0.690 ± 0.002 | 0.627 | N/A | none |
| 4 | 0.708 ± 0.002 | 0.626 | N/A | none |
| 3 | 0.690 ± 0.001 | 0.636 | N/A | none |
| 2 | 0.685 ± 0.001 | 0.613 | N/A | none |

Table 2. Challenge scores for our domain-invariant model (trained on all public training data) using 5-fold cross validation on the public training set and repeated scoring on the hidden validation set. There are no *Test* results for this model due to a technical error.

| Leads | Training | Validation | Test | Ranking |
|-------|--------------------|------------|------|---------|
| 12 | 0.370 ± 0.0236 | 0.650 | 0.42 | none |
| 6 | 0.398 ± 0.0448 | 0.578 | 0.42 | none |
| 4 | 0.427 ± 0.0156 | 0.634 | 0.42 | none |
| 3 | 0.442 ± 0.0233 | 0.645 | 0.44 | none |
| 2 | 0.410 ± 0.0377 | 0.619 | 0.38 | none |

Table 3. Challenge scores for our baseline model (trained on CPSC, CPSC-Extra and G12EC) using 5-fold cross validation on the public training set, repeated scoring on the hidden validation set, and one-time scoring on the hidden test set. We did not receive a rank for this model on the hidden test set.

tract domain-invariant features from 12-lead or reduced-lead ECGs (Table 1). For the baseline model, training on CPSC, CPSC-Extra and G12EC (same domain to the online validation data) leads to a relatively low score for all local training data but a high score for online validation data (Table 3).

Although the domain-invariant model showed improvement on the unseen domain when compared to the baseline model, both models had a noticeable drop in performance for the unseen domain. This may indicate that domain-invariant features are only part of the solution for making a more general model, as there may be genuine differences such as different diagnostic criteria which cannot be accounted for with domain-invariant features alone. In future work, we will apply this model to the test set, and systematically optimise hyperparameters.

References

[1] Reyna MA, Sadr N, Perez Alday EA, Gu A, Shah A, Robichaux C, et al. Will Two Do? Varying Dimensions in Electrocardiography: the PhysioNet/Computing in Cardiology Challenge 2021. *Computing in Cardiology* 2021;48:1–4.

[2] Alday EAP, Gu A, Shah AJ, Robichaux C, Wong AKI, Liu C, et al. Classification of 12-lead ECGs: the PhysioNet/Computing in Cardiology Challenge 2020. *Physiological Measurement* jan 2021;41(12):124003.

[3] Zhao Z, Fang H, Relton SD, Yan R, Liu Y, Li Z, et al. Adaptive Lead Weighted ResNet Trained with Different Duration Signals for Classifying 12-lead ECGs. In *2020 Computing in Cardiology*. IEEE, 2020; 1–4.

[4] Hu J, Shen L, Sun G. Squeeze-and-Excitation Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2018; 7132–7141.

[5] Liu F, Liu C, Zhao L, Zhang X, Wu X, Xu X, et al. An Open Access Database for Evaluating the Algorithms of Electrocardiogram Rhythm and Morphology Abnormality Detection. *Journal of Medical Imaging and Health Informatics* 2018;8(7):1368—1373.

[6] Tihonenko V, Khaustov A, Ivanov S, Rivin A, Yakushenko E. St Petersburg INCART 12-lead Arrhythmia Database. *PhysioBank PhysioToolkit and PhysioNet* 2008;Doi: 10.13026/C2V88N.

[7] Boussejot R, Kreiseler D, Schnabel A. Nutzung der EKG-Signaldatenbank CARDIODAT der PTB über das Internet. *Biomedizinische Technik* 1995;40(S1):317–318.

[8] Wagner P, Strodthoff N, Boussejot RD, Kreiseler D, Lunze FI, Samek W, et al. PTB-XL, a Large Publicly Available Electrocardiography Dataset. *Scientific Data* 2020;7(1):1–15.

[9] Zheng J, Zhang J, Danioko S, Yao H, Guo H, Rakovski C. A 12-lead Electrocardiogram Database for Arrhythmia Research Covering More Than 10,000 Patients. *Scientific Data* 2020;7(48):1–8.

[10] Zheng J, Cui H, Struppa D, Zhang J, Yacoub SM, El-Askary H, et al. Optimal Multi-Stage Arrhythmia Classification Approach. *Scientific Data* 2020;10(2898):1–17.

[11] Martis RJ, Acharya UR, Min LC. ECG beat classification using PCA, LDA, ICA and Discrete Wavelet Transform. *Biomedical Signal Processing and Control* 2013;8(5):437–448.

[12] Engelse WA, Zeelenberg C. A single scan algorithm for QRS-detection and feature extraction. *Computers in Cardiology* 1979;6:37–42.

[13] Zhang W, Peng G, Li C, Chen Y, Zhang Z. A New Deep Learning Model for Fault Diagnosis with Good Anti-Noise and Domain Adaptation Ability on Raw Vibration Signals. *Sensors* 2017;17(2):425.

[14] He K, Zhang X, Ren S, Sun J. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2016; 770–778.

[15] Ganin Y, Lempitsky V. Unsupervised Domain Adaptation by Backpropagation. In *International Conference on Machine Learning*. PMLR, 2015; 1180–1189.

Address for correspondence:

Zuogang Shang

International center of Machinery, School of Mechanical Engineering, Xi'an Jiaotong University, Xi'an, China.
administrator@stu.xjtu.edu.cn