SmartHeart: Deep Learning-Based Analysis of 12-Lead Electrocardiograms for Cardiovascular Disease Risk Prediction

Evan Wen, JP Salvatore, Alan Zhong, Dr. Jolly

Pingry Research

Research Week 2022: SmartHeart IRT

Introduction

In the US, Heart Disease is the leading cause of death for both men and women, causing about 1 in every 4 deaths. The most common method to diagnose Heart Disease is by obtaining a cardiovascular procedure called an electrocardiogram (ECG), which consists of placing electrodes on various parts of the body to form a graph of the heart's electrical activity.
Because of limitations in trained cardiologists and existing deep learning models, we're seeking to answer the following question:

How can state-of-the-art deep learning be leveraged for automated analysis of 12-lead ECGs?

We're currently working on accurately classifying ECGs as either healthy or with 1 of 27 Heart Diseases. We source our ECGs from online datasets, and we also obtain real, paper ECGs from a local cardiologist.

Methods

We've tried 3 different machine-learning based approaches to this problem. Before training the model, we first processed the ECG as a 12 by 7500 grid, derived from 12 leads (signals) and 7500 (500Hz for 15 seconds) sampled points on each lead.

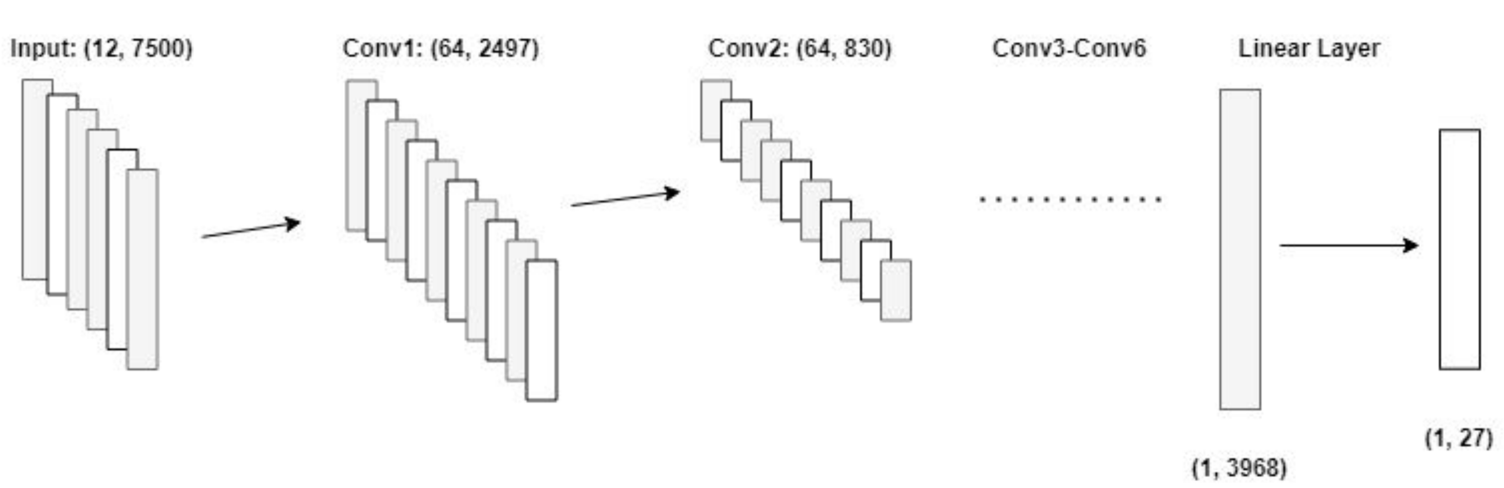

Firstly, we used a standard convolutional neural network (CNN), which runs 1D convolutions across our input ECG to eventually achieve an output layer that classifies the ECG into 1 of 27 diseases.

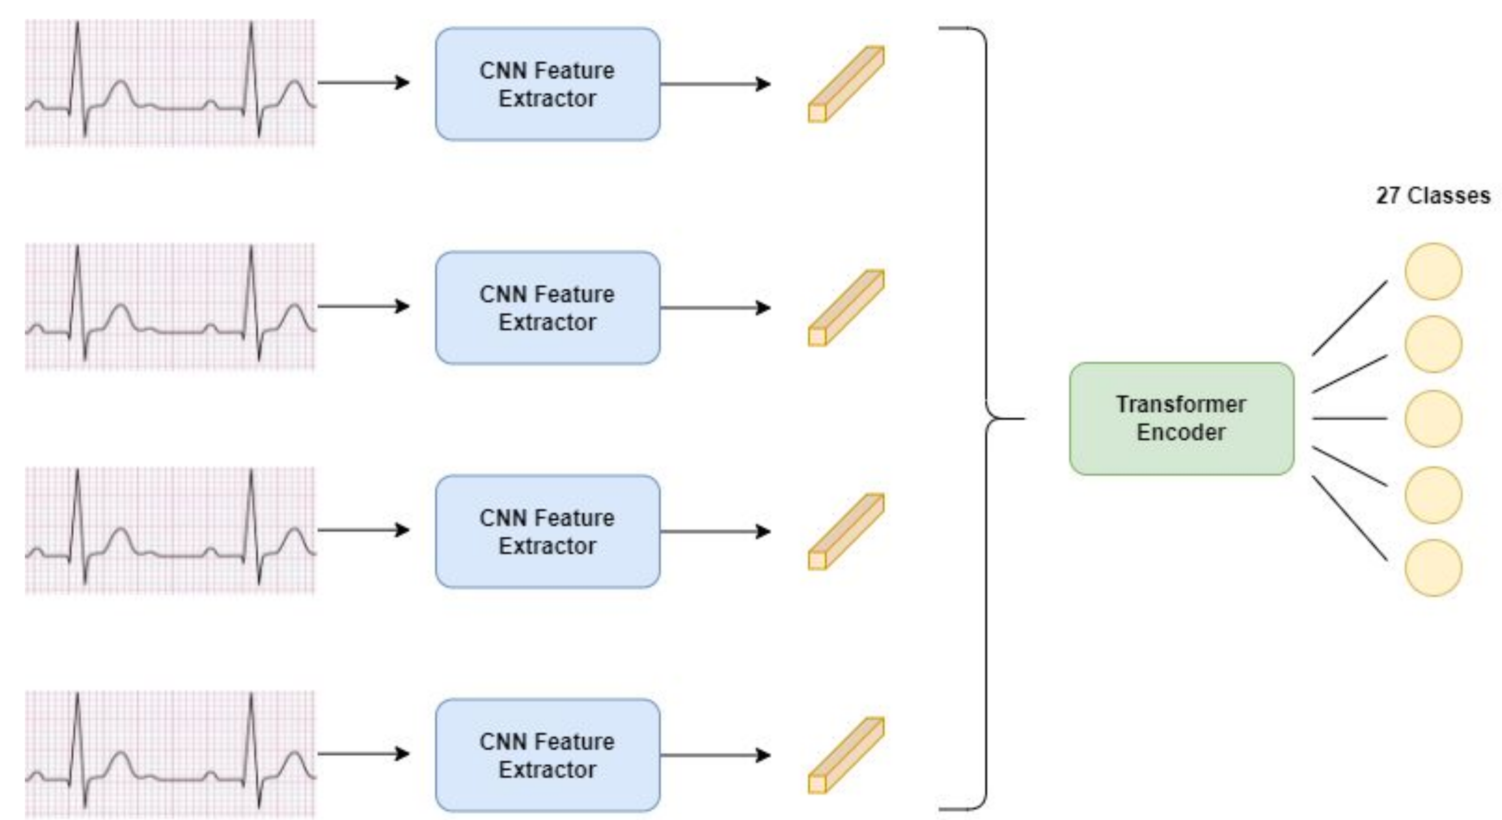

Secondly, we tried dividing each ECG into 5 segments of equal length, which are then passed through CNNs to extract a feature vector. These vectors are then passed through a transformer encoder to achieve our final output classification.

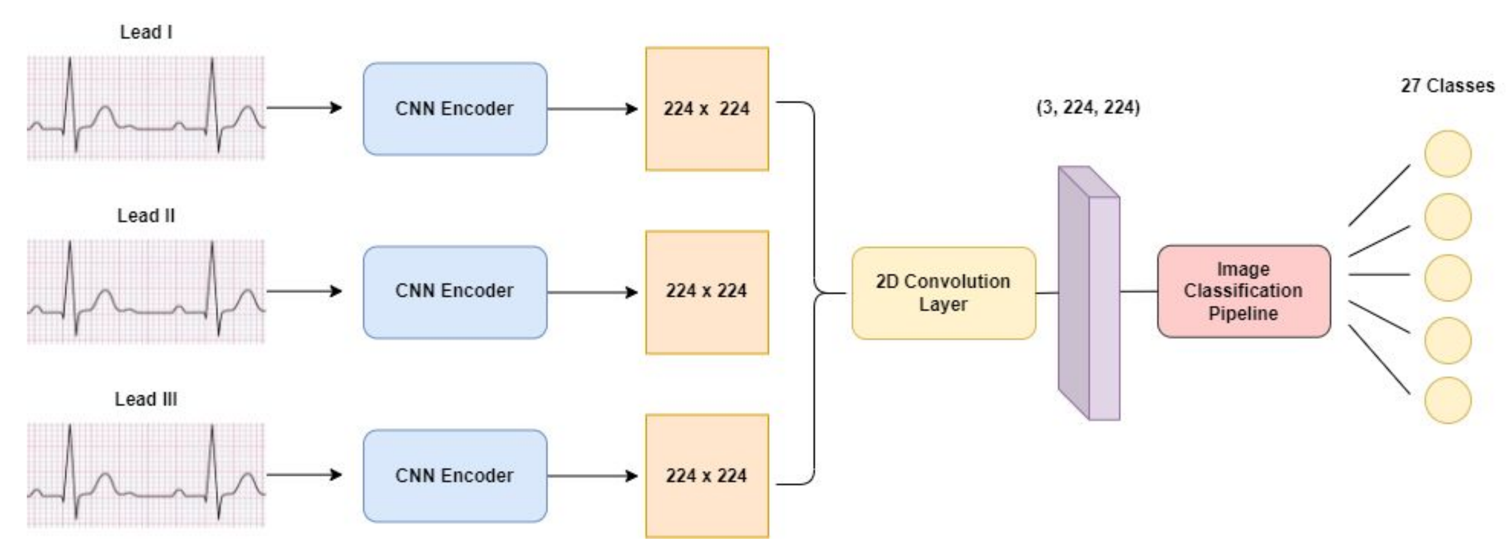

Finally, we applied image classification pipelines on image representations of ECGs. We use CNN encoders to transform each lead into a 2D array, before concatenating them together. We then pass the leads into a 2D convolution to obtain an RGB image representation, and we classify the image into our final output.

Results

Approach 1

Feature Extraction	Head	Accuracy (%)
1D CNN	None Transformer	75.84 ± 2.04 79.98 ± 0.92
1D ResNet	None **Transformer**	80.14 ± 0.92 **81.43 ± 0.49**

Approach 2

# of Splits	Data Augmentation?	Training Accuracy (%)	Validation Accuracy (%)
3	No **Yes**	94.36 ± 2.75 72.90 ± 3.18	43.67 ± 1.54 **53.78 ± 0.86**
5	No Yes	92.62 ± 1.06 93.75 ± 3.24	41.59 ± 0.80 49.84 ± 1.19

Approach 3

Pipeline	Data Augmentation?	Training Accuracy (%)	Validation Accuracy (%)
ResNet-50	No Yes	94.82 ± 1.72 91.28 ± 1.89	49.88 ± 1.64 43.51 ± 2.14
DenseNet-121	No Yes	82.91 ± 1.60 90.14 ± 1.82	50.91 ± 1.52 43.49 ± 1.48

Future Directions

In the future, we want to work with hybrid ECGs, created by adding anomalous regions from a diseased ECG onto a healthy ECG. We hypothesize that using these hybrid ECGs, we'd be able to forecast a patient's risk for later diagnosis of a heart disease, which we've been aiming to do.

Our goals for using hybrid ECGs are to evaluate the accuracy of generated hybrid ECGs, and to apply these generated hybrid ECGs for risk prediction. So far, we've developed a theoretical implementation of generating hybrid ECGs and evaluating their accuracy.

References

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., … Polosukhin, I. (2017). *Attention Is All You Need*. doi:10.48550/ARXIV.1706.03762

He, K., Zhang, X., Ren, S., & Sun, J. (2015). *Deep Residual Learning for Image Recognition*. doi:10.48550/ARXIV.1512.03385

Huang, G., Liu, Z., van der Maaten, L., & Weinberger, K. Q. (2016). *Densely Connected Convolutional Networks*. doi:10.48550/ARXIV.1608.06993

Park, T., Zhu, J.-Y., Wang, O., Lu, J., Shechtman, E., Efros, A. A., & Zhang, R. (2020). *Swapping Autoencoder for Deep Image Manipulation*. doi:10.48550/ARXIV.2007.00653

Acknowledgements: Dr. Morgan D'Ausilio, Dr. Azadeh Samadani, Dr. Robert Masci