Using lightweight image classifiers for electrocardiogram classification on embedded devices

Niek Heinen
University of Twente
P.O. Box 217, 7500AE Enschede
The Netherlands
n.l.heinen@student.utwente.nl

ABSTRACT

A large part of deep learning research is devoted to image classification. The research in this paper will show that the same neural networks developed for image classification can also be used to accurately classify electrocardiograms (ECGs). Even though this is a novel approach for ECG classification, the early results appear to be promising. Since both accuracy and efficiency are much valued in the applications of ECG classifiers the research will focus on using lightweight image classifiers such as ResNet and MobileNet. In order to (maximally) utilise the architecture of the image classifiers, we need to cleverly reshape the ECG signals. In this paper, we will investigate numerous ways of doing so. Using the VEB en SVEB evaluation metrics, the research in this paper will be compared to the state-of-the-art. The best performing ECG classifier presented in this paper achieved a VEB- and SVEB-accuracy of 96.8% and 98.1% respectively.

Keywords

ECG classification, Deep Learning, ResNet, MobileNet, AAMI, VEB and SVEB evaluation metrics, Embedded Devices.

1. INTRODUCTION

An electrocardiogram (ECG) is a measurement of electrical activity in the heart. An ECG is usually represented as a graph with voltage on the y-axis versus time on the x-axis, figure 1 is an example of such a graph. All over the world ECGs are used to detect numerous kinds of cardiac abnormalities. As the leading leading cause of mortality worldwide [22], cardiovascular diseases can be better treated and earlier identified with the help of an automated ECG classifier.

Many ECG classifying software solutions exist today, most of these solutions are machine learning based classifiers that can precisely detect some cardiac abnormalities [15]. Be that as it may, no cardiologist-replacing all-purpose ECG classifier exist yet. However, with the rise of new promising deep learning applications in the healthcare sector [3], this might change. The best performing state-of-the-art ECG classifiers mostly make use of denoising autoencoders (DAE) in combination with a deep neural network [20, 23], or with the use of support vector machines (SVM) [17, 2]. Within the the ECG classifying research community there are no clear guidelines on what classes should be classified. Many researchers use different classes which makes it difficult to fairly compare most ECG Classifiers. The variation in focus and classes in different research is most likely caused by the large number of possible cardiac disorders that could be derived from an ECG.

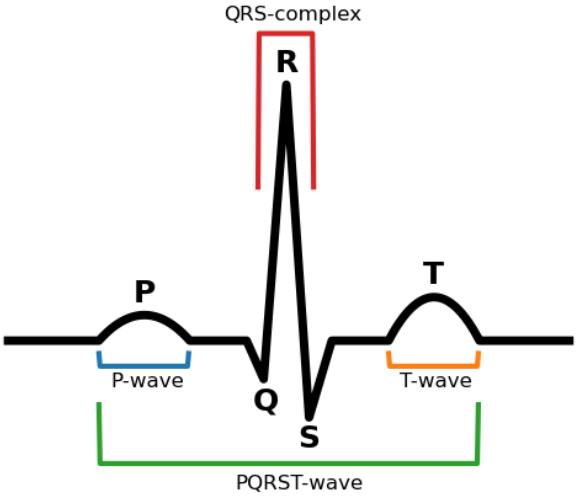

Figure 1. ECG of a regular cardiac cycle.

In this paper a novel approach for classifying ECGs will be presented. The novel approach involves the use of image classifiers, by doing so, all image classification research can now be used to classify ECGs. As all images have height, width and colour image classifiers are made for three dimensional input. ECG signals are one dimensional, so in order to make use of an image classifiers architecture, the ECG signal needs to be reshaped. In this research we will investigate numerous ways of doing so. As many useful applications of the ECG classifier are in wearable devices, it is not only important that the classifier is accurate, but also efficient. Some image classifiers can be rather complex networks, so in order to assure the efficiency of the ECG classifier, the research will focus on some of the more efficient and lightweight image classifiers. More specifically, ResNet50 and ResNet50V2 [8], DenseNet121, 169 and 201 [10], NasNetMobile [25] and, MobileNet and MobileNetV2 [9].

33rd Twente Student Conference on IT July 3rd, 2020, Enschede, The Netherlands.

2. RELATED WORK

Rahhal *et al.* [20], observed that the feature representation of ECG signals in most state-of-the-art methods on relied handcrafted features. They found an automatic way of doing this using an DAE.

The reseach of Xia *et al.* [23] achieved very accurate results[1] on multiple public ECG databases with the use of stacked DAEs combined with a deep neural network.

Ma *et al.* [16] popularised the use of SVMs for ECG classification. With the use of particle swarm optimisation, they achieved an overall accuracy of 89.72%.

Rajesh *et al.* further improved the performance of SVM based ECG classifiers. They also classified more specific beats, meaning they had to more classes compared to other state-of-the-art research. They achieved an overall accuracy of 94.12%.

In the work of Acharya *et al.* [4] and Zheng **et al.** [24] it has been found that convolutional neural networks can also be used to effectively classify ECGs.

Le Cun *et al.* [1] are often credited for developing the first real deep convolutional neural network. Krizhevsky et al. invented AlexNet [14] and made great progress in creating more accurate and faster CNNs for image recognition.

Inspired by constructs of neurons in the human brain, residual neural networks (ResNets) were invented by He *et al.* in 2016 [8]. ResNets drastically improved the training times of deep neural networks while maintaining accurate results. ResNets achieve this by making "shortcuts" between between layers which creates a much smaller (and therefore easier to train) neural network.

3. ECG CLASSIFICATION

3.1 Classes

Classifying ECGs can be quite a challenging task as there are a great number of different types of classes we could classify beats as. Even experts sometimes disagree on the classification of certain beats as the differences between classes is not always obvious. According to the Association for the Advancement of Medical Instrumentation (AAMI) standard there are 5 important classes we should classify ECGs as, Normal (N), Supraventricular (S), Ventricular (V), Fusion of normal and ventricular (often referred to as Fusion beat or F) and Unknown beats (Q) [5]. A visual representation of of every AAMI class is presented in figure 2, every class is presented as a traditional one dimensional signal, and its reshaped version. In section 4.1 a more elaborate explanation of the reshaped signal is given.

3.2 Data set

The data set used to train the ECG classifier will be the MIT-BIH Arrhythmia Database [18]. This is a publicly available database containing 48 records all of which are 30 minutes long. They are 2 leads ECGs recorded at 360Hz. The ECGs were obtained from 47 subjects between 1975 and 1979. The 48 records are divided into two categories, the first 23 belong to the "100 series", these records are representative of general population. The remaining 25 belong to the "200 series", these contain uncommon but clinically important cardiac arrhythmias that would not be well represented in a small random sample [19]. Multiple independent experts have annotated every beat, in table 1 you can find an overview all beats annotations (BA) contained in the database, their frequency, and their corresponding AAMI class.

[1]The exact results of Xia *et al.*'s method can be found in table 5

2(a) Example of Normal beat, BA class N

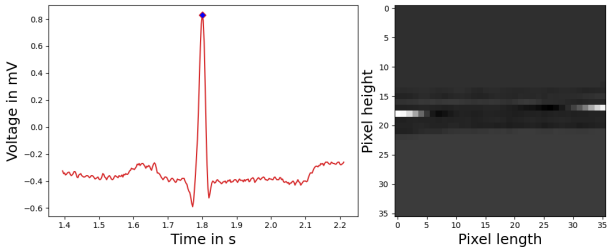

2(b) Example of Supraventricular beat, BA class A

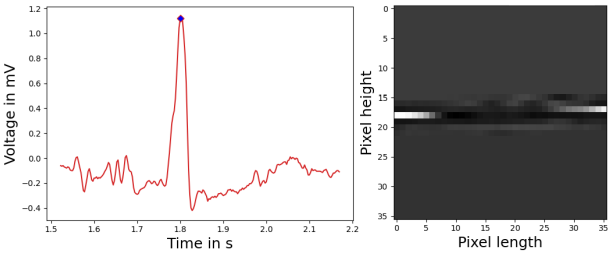

2(c) Example of Ventricular beat, BA class V

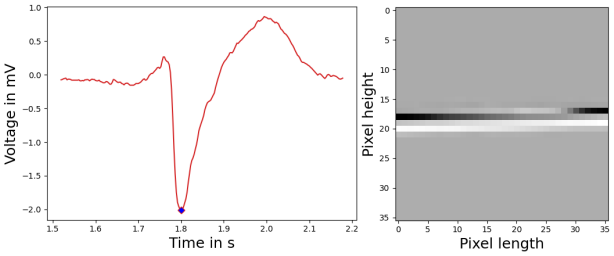

2(d) Example of Fusion beat, BA class F

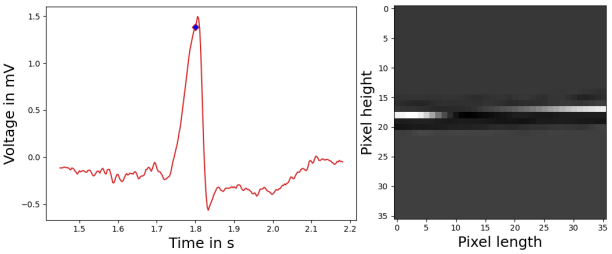

2(e) Example of Unknown beat, BA class Q

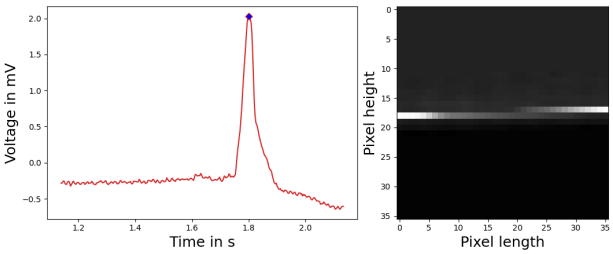

Figure 2. Examples of all AAMI classes as both a traditional one dimensional signal, and a visual representation of the reshaped signal.

The MIT-BIH database is well-known among researchers occupied with creating ECG classifiers, [6, 13, 16, 20, 23] all used the database to train and analyse the accuracy of their developed solutions. The data will be split into the following two sets; $DS1$ = {101, 106, 108, 109, 112, 114, 115, 116, 118, 119, 122, 124, 201, 203, 205, 207, 208, 209, 215, 220, 223, 230} and $DS2$ ={100, 103, 105, 111, 113, 117, 121, 123, 200, 202, 210, 212, 213, 214, 219, 221, 222, 228, 231, 232, 233, 234}. $DS1$ will be used to train the classifier, $DS2$ will be used to test the performance. The

AAMI	BA	Description	Frequency	Count
Normal	N	Normal beat	70.86%	74790
	L	Left bundle branch block beat	7.65%	8075
	R	Right bundle branch block beat	6.88%	7259
	e	Atrial escape beat	0.02%	16
	j	Nodal (junctional) escape beat	0.22%	229
Supraventricular	A	Atrial premature beat	2.41%	2546
	a	Aberrated atrial premature beat	0.14%	150
	S	Supraventricular premature beat	0.01%	2
	J	Nodal (junctional) premature beat	0.08%	83
Ventricular	V	Premature ventricular contraction	6.75%	7124
	!	Ventricular flutter wave	0.45%	472
	E	Ventricular escape beat	0.10%	106
Fusion Beat	F	Fusion of ventricular and normal beat	0.76%	803
Unkown beats	/	Paced beat	3.43%	3620
	f	Fusion of paced and normal beat	0.25%	260
	Q	Unclassifiable beat	0.01%	15

Table 1. Overview of MIT-BIH Database beat annotations.

frequency of the classes is proportionate in these datasets, they are the same datasets used in the research of [6, 20]. The records with serial number 102, 104, 107 and 217 are in neither of the datasets, most ECG classifying researchers exclude these records as they contain paced heartbeats. In order to make fair comparisons in the final evaluation of the classifier, these records will also be excluded in this research.

4. METHODOLOGY

4.1 Signal reshaping

To make use of the image classifiers we need need to reshape our ECG signal. As aforementioned, all modern image classifiers take three dimensional tensors as input. The minimal dimensions of the input required for most image classifiers is $32 \times 32 \times 3$, the first two numbers simply mean that the image has to have at least a height and width of 32 pixels. The third and final dimension represents the three RGB channels, i.e. every pixel has three values that make up its colour. Since an ECG signal is one dimensional, we need to reshape the signal in order to make it work with image classifiers' architecture. Firstly, a good window size has to be determined, which can be more difficult than one might expect. This is mostly because of the large differences in cardiac cycles. E.g. some patient's average heartbeat lasts only $600ms$, while for other patients it lasts $1500ms$. And then of course there are the outliers, some "beats" in the database last for over $5s$. In this research we have chosen for a window size of $3.6s$. Since the ECGs are recorded at $360Hz$, this means we have $3.6 \times 360 = 1296$ measurements per window, which is the perfect square of 36. We can then reshape the signal into a 36×36 matrix, and copy it into all three RGB channels, similar to how grayscale images are passed into image classifiers. The QRS-complex is the most significant part of the ECG signal and therefore also the most important for classifying beats [13]. For this reason, the top of the QRS-complex is always placed in the centre of the window (in visual representations of the signal, a blue diamond is placed at the top of the QRS-complex for clarifying purposes).

3(a) Signal, and tensor format *OS*

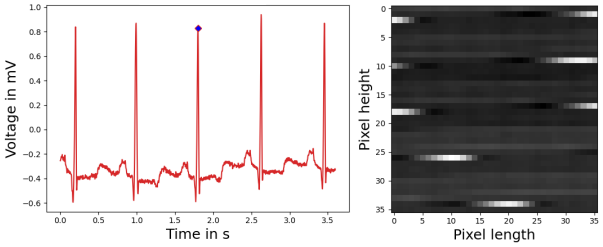

3(b) Signal, and tensor format *IS*

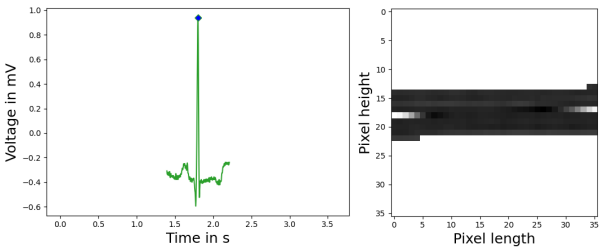

3(c) Signal, and tensor format *ES*

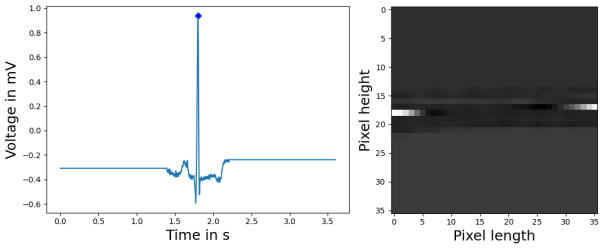

3(d) Signal, and tensor format *MS*

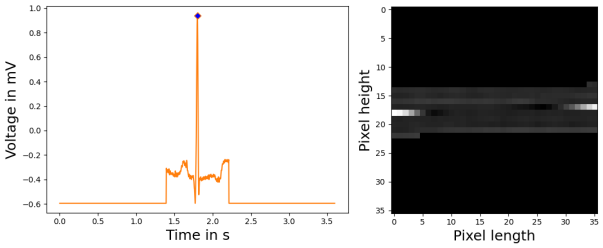

Figure 3. Visual representation of the different ECG formats, as signal and as tensor.

With a 3.6s window, it often occurs that multiple other beats are contained within the same window, because this might have a negative effect on the classifiers' performance, four different signal formats are presented; the original signal *OS* where no modifications to the signal are made, the isolated signal *IS* where only the centre heartbeat's PQRST-wave is represented in the window, the extended signal *ES* where the first and last value of the centre heartbeat's PQRST-wave are extended to the beginning and end of the window, and finally, the minimised signal *MS* where all values that are not in the centre PQRST-wave are minimised. Since these four formats result in different tensors when reshaped, all are interpreted differently by the image classifier. In this research, we will examine what the most efficient format is. In figure 3[2] a visual representation of the different formats can be found.

4.2 Classifier architecture

The ECG classifier has three core components; the signal reshaper, the image classifier, and the interpretation layer. In figure 3 you can find an schematic overview of the ECG classifier's architecture. The first core component of the ECG classifier is the signal reshaper, this component is responsible for transforming the one dimensional ECG signal into into something the image classifier can work with. To the determine the best performing signal shape, the image classifier's performance with different signal formats will be investigated. The largest and second core component is the image classifier. The image classifier could be considered the "engine" of the ECG classifier, as the ECG classifier is highly dependant its engines performance. However, this engine can easily be replaced by any other image classifier. This allows for easy testing between the accuracy of the different image classifiers. In order to examine what kind of image classifier's architecture is best suited for serving as part of the ECG classifier, we will compare the performance of multiple image classiers. When making use of the image classifiers, we can either use the publicly available pre-trained weights used in the ImageNet [21] competition, or have the weights be initialised randomly. In table 2 an overview of all researched image classifiers is shown, as well as their ImageNet top 1 accuracy, number of default trainable parameters (*DTP*), number of layers, and the size of the image classifier in MB. Generally, the more trainable parameters a neural network has, the longer it takes to train and run the classifier. In order to use the image classifier as an ECG classifier, an interpretation layer (*IL*) is required. This is the third and final core component of the ECG classifier. In order to determine what the best kind of structure for the *IL* is, three interpretation layer configurations *ILCs* are proposed; *ILC1*: one fully connected layer with five neurons. *ICL2*: two fully connected layers with 64 and 5 neurons respectively, and *ICL3*: three fully connected layers with 124, 64 and 5 neurons respectively. In this research we will investigate the optimal configuration of these interpretation layers.

4.3 Evaluation

To evaluate the ECG classifier we will use the VEB, and SVEB evaluation metrics, as introduced by [7]. The Ventricular and Supraventricular are the classes that come with the most health risks, and are therefore the most important to identify. The VEB and SVEB metrics are respectively Ventricular versus all other classes, and Supraventricular versus all other classes. These are well-known evaluation metrics in ECG classification research, and also

[2]The ECG used in this figure has *BA*-classification N and is the 200th beat in record 100 of the MIT-BIH database.

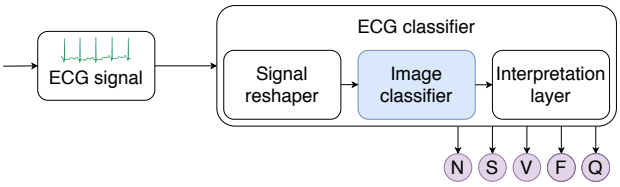

Figure 4. Schematic summary of the ECG classifier's architecture.

Image Classifier	Acc	*DTP*	Layers	MB
ResNet50	0.749	23,534,592	172	98
ResNet50V2	0.760	23,519,360	191	98
DenseNet121	0.750	6,953,856	428	33
DenseNet169	0.762	12,484,480	596	57
DenseNet201	**0.773**	18,092,928	708	80
NasNetMobile	0.744	4,232,978	771	23
MobileNet	0.704	3,206,976	88	16
MobileNetV2	0.713	**2,223,872**	156	**14**

Table 2. All researched image classifiers, their accuracy in the ImageNet competition and some information about their architecture.

used by [23]. When comparing the classifier's performance, the following VEB and SVEB metrics are used; accuracy (Acc), positive predictivity (Pp), Sensitivity (Se) and Specificity (Sp). Since the same metrics are used in some of the state-of-the-art related research, the ECG classifier from this paper can be fairly compared to other classifiers.

4.4 Open source

Once the research has been published, all code written for the project will be made publicly available on GitHub. There, everyone is free to use the code for further research, or improve upon it.

5. RESULTS

5.1 Setup

Because of the limited amount of time available for this research, the optimal configurations for each of the classifier's components were determined by testing them with only some of the other component's configuration. E.g. when investigating the best performing interpretation layer configuration, only MobileNetV2 and ResNetV2 were considered. Even though ResNetV2 performed better with *ILC1*, it is possible that DenseNet121 or NasNetMobile is more more accurate when combined with *ILC3*.

First, the best performing architecture configuration was determined. Table 3 shows an overview of different architecture configurations and their corresponding performance. The architecture configuration includes (From left to right as shown in table 3): *ILC* , the interpretation layer configuration that was used. Trained layers, the layers that were trained in the model, e.g. *DTP + IL* means that all default trainable parameters and the interpretation layers were trained. Epochs, the number of epochs the model has trained for. The epoch steps are 30, 60, 100 and 150. If after an epoch step the model does not seem to be improving any further, no more epoch steps are taken. The fourth configuration column s/epoch represents the average number of seconds it took to execute one epoch[3]. Weights, when using the image classifiers it

[3]The models were all trained in Google Colab notebooks

Image Classifier	Configurations					Accuracy			
	ILC	Trained layers	Epochs	s/epoch	Weights	Test	VEB	SVEB	Train
MobileNetV2	*ILC1*	*IL*	30	9	ImageNet	0.872	0.935	0.958	0.967
MobileNetV2	*ILC1*	*IL*	60	9	ImageNet	0.834	0.918	0.943	0.971
MobileNetV2	*ILC1*	*IL*	100	9	ImageNet	0.814	0.893	0.938	0.975
MobileNetV2	*ILC2*	*IL*	30	10	ImageNet	0.883	0.953	0.947	0.995
MobileNetV2	*ILC2*	*IL*	60	10	ImageNet	0.798	0.915	0.912	0.997
MobileNetV2	*ILC3*	*IL*	30	10	ImageNet	0.845	0.924	0.942	0.991
MobileNetV2	*ILC3*	*IL*	60	10	ImageNet	0.834	0.935	0.925	0.966
MobileNetV2	*ILC1*	*DTP + IL*	30	52	ImageNet	0.802	0.897	0.834	0.996
MobileNetV2	*ILC1*	*DTP + IL*	60	52	ImageNet	0.816	0.904	0.937	0.996
MobileNetV2	*ILC1*	*DTP + IL*	100	52	ImageNet	**0.904**	**0.950**	**0.962**	0.999
MobileNetV2	*ILC1*	*DTP + IL*	150	52	ImageNet	0.867	0.968	0.952	0.999
ResNetV2	*ILC1*	*IL*	30	38	ImageNet	**0.923**	**0.968**	**0.963**	0.949
ResNetV2	*ILC1*	*IL*	60	38	ImageNet	0.917	0.908	0.956	0.960
ResNetV2	*ILC1*	*IL*	100	38	ImageNet	0.803	0.894	0.947	0.968
ResNetV2	*ILC2*	*IL*	30	38	ImageNet	0.903	0.950	0.963	0.949
ResNetV2	*ILC2*	*IL*	60	38	ImageNet	0.914	0.957	0.963	0.960
ResNetV2	*ILC2*	*IL*	100	38	ImageNet	0.913	0.957	0.963	0.955
ResNetV2	*ILC3*	*IL*	30	38	ImageNet	0.875	0.921	0.963	0.955
ResNetV2	*ILC3*	*IL*	60	38	ImageNet	0.899	0.943	0.963	0.955
ResNetV2	*ILC3*	*IL*	100	38	ImageNet	0.913	0.957	0.963	0.955
ResNetV2	*ILC1*	*DTP + IL*	30	152	ImageNet	0.826	0.897	0.953	0.992
ResNetV2	*ILC1*	*DTP + IL*	60	152	ImageNet	0.801	0.875	0.930	0.999
ResNetV2	*ILC1*	*DTP + IL*	100	152	ImageNet	0.792	0.860	0.921	0.999

Table 3. Classifier performance with different architecture configurations.

is possible to either use the pre-trained ImageNet weights, or have all weight be randomly initialised. Since, a lucky random initialisation can have a big impact on the classifiers performance, numerous attempts of initialising the weight randomly were made, unfortunately without any success, as the classifiers performance with the ImageNet weights were in every scenario considerably better. Table 3 contains the highest performances obtained with the specified configurations.

Secondly, the optimal image classifier was determined by applying the best performing architecture configuration to all image classifiers, and choosing the one with the best overall results.

Finally, the most efficient signal shape was determined. Because of time constraints, the best performing architecture and image classifier were determined while only using signal format *ES*. This is because it is presumably the optimal signal shape, as the QRS-complex is the most significant feature in this tensor, whereas formats *IS* and *MS* also have high contrast between the the PQRST-wave and the rest of the signal. The best performing architecture and image classifier are then combined with the four signal formats, to determine the optimal signal format.

5.2 Experimental results

In table 3, the experimental results for determining the optimal architecture are presented. Even though MobileNetV2's optimal architecture configuration differs from ResNet50V2's configuration, the following optimal architecture configuration was determined. The best configuration is *ILC1*, training only the interpretation layer for 30 epochs, and initialise the weights as the pre-trained ImageNet weights. The choice for this configuration was made because of two reasons; these results are better when compared to the optimal configuration for MobileNetV2, and simply because it would take significantly longer to train all other image classifiers with the optimal MobileNetV2 configuration. E.g. training both the *DTP* and *IL* for 100 epochs took over four hours when applied to ResNet50V2. In the rest of the results section, all models will have these architecture configurations applied.

In table 4, the performances of all the different image classifiers are presented. All models are trained for 30 epochs with *ILC1* and signal format *ES*. With these configurations, ResNet50V2 appears to be performing the best. Finally, the optimal signal format was determined, from the results presented in table 5 we can derive that *ES* is the optimal format. However, there could be argued against this conclusion, as every other optimal component was determined by using the *ES* format as a baseline, and it therefore has an unfair advantage over the other formats.

6. DISCUSSION

Considering the likelihood that there exists an even more optimal configuration, combined with the rather decent accuracy of the best performing classifier presented in this paper, this novel approach of classifying ECGs suggests potential. This way of using image classifiers, as the "engine" of an ECG classifier is an interesting way to make use of the state-of-the-art image classification research. If an even better image classifier is developed and published, it

with GPU hardware acceleration enabled. These number might be difficult to reproduce, since every time you reconnect to a Google Colab Notebook, you could be assigned different hardware.

Image Classifier	Acc	VEB				SVEB			
	(General)	Acc	Pp	Se	Sp	Acc	Pp	Se	Sp
ResNet50	0.901	0.952	0.614	0.692	**0.970**	0.963	0.000	0.000	**1.000**
ResNet50V2	**0.924**	**0.968**	**0.679**	0.586	0.963	**0.981**	0.213	**0.069**	0.990
DenseNet121	0.759	0.863	0.307	0.875	0.863	0.963	**0.517**	0.017	0.999
DenseNet169	0.790	0.901	0.389	**0.911**	0.901	0.962	0.185	0.008	0.999
DenseNet201	0.883	0.958	0.621	0.899	0.962	0.958	0.103	0.002	0.994
NasNetMobile	0.854	0.928	0.488	0.839	0.920	0.963	0.184	0.057	0.999
MobileNet	0.835	0.866	0.111	0.153	0.915	0.963	0.204	0.001	0.999
MobileNetV2	0.872	0.935	0.470	0.772	0.940	0.958	0.092	0.023	0.991

Table 4. Image classifier comparison.

Image classifier	Format	VEB				SVEB			
		Acc	Pp	Se	Sp	Acc	Pp	Se	Sp
ResNet50V2	*OS*	0.842	0.394	0.773	0.842	0.921	0.127	0.004	0.952
ResNet50V2	*IS*	0.892	0.425	**0.869**	0.931	0.963	0.002	0.004	**0.999**
ResNet50V2	*ES*	**0.968**	**0.679**	0.586	**0.963**	**0.981**	**0.213**	**0.069**	0.991
ResNet50V2	*MS*	0.942	0.643	0.531	0.942	0.952	0.120	0.012	0.982

Table 5. Signal format performance comparison.

Method	VEB				SVEB			
	Acc	Pp	Se	Sp	Acc	Pp	Se	Sp
Ince et al. [11]	0.976	0.874	0.834	0.981	0.961	0.567	0.621	0.985
Kiranyaz et al. [12]	0.986	0.895	0.950	0.981	0.964	0.621	0.646	0.986
Xia et al. [23]	**0.995**	**0.973**	**0.979**	**0.997**	**0.997**	**0.956**	**0.961**	**0.999**
This paper	0.968	0.679	0.586	0.963	0.981	0.213	0.069	0.991

Table 6. Comparison to the state-of-the-art research.

can without much effort be applied in the ECG classifier. From the results in tables 3, 4 and 5, we can also conclude that the ECG classifier seems to particularly struggle with the Supraventricular class. As seen in the visualisation of the tensors in figure 2, the current setup does not show a very big difference between the Normal and Supraventricular class. Experimenting with different signal formats, even different from the ones presented in this paper, could help the ECG classifier perform better on the Supraventricular class. There are also disadvantages to this approach. In healthcare it is of high importance that a diagnosis is not only made, but can also be explained. At this moment, the classifier could be considered a black-box. We can see the classification it makes, but we have no idea how it reached this classification. There is plenty of research being done as how making deep neural networks explainable. If ever applied in the real world, it would be necessary for the ECG classifier to be able to point to the features in the original signal that led to the decision of giving it a certain classification.

There are numerous ways to further improve upon the accuracy of the ECG classifier. Right now the raw ECG signal is fed to the signal reshaper, however, research in [23] has shown denoising the signal could help with feature extraction and improve the overall performance of the classifier. Also barely any hyper-parameter optimisation was preformed. Exploring different optimisers and loss functions could drastically improve the classifiers' performance. The classifier presented in this paper has a low Pp and Se, these results could be enhanced with improved feature extraction, or by having the classifier train or more data. Besides the regular MIT-BIH arrhythmia database as used in this paper, there also exists an MIT-BIH Supraventricular Arrhythmia Database. This database contains a larger portion of Supraventricular beats. This database can be used to further enhance the classifier's performance. Plenty of research is still to be done in the optimal signal formatting. Especially considering 12-lead ECGs, the RGB channels could used for holding different leads, rather than duplicating one lead. Another approach for the signal reshaping could be to turn the one dimensional ECG signal into a spectrogram. This could perhaps help improve feature extraction.

7. CONCLUSION

In this paper a novel approach for classifying ECGs is presented. The novel approach makes use use of image classifiers, this means the state-of-the-art image classification research can now be applied for ECG classifying research. The early results presented in this paper are worse when compared to SVM or DAE based ECG classifiers from related works. However, when taken into consideration the minimal optimisation that was performed, these early results suggest potential. The biggest shortcoming of the current presented solution is its struggle with Supraventricular class, this could be improved with enhanced feature extraction, better signal reshaping, or more training data. In order to determine if this novel approach is of any value, further research will have to be conducted.

8. REFERENCES

[1] Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2323, 1998.

[2] Classification of ECG heartbeats using nonlinear decomposition methods and support vector machine. *Computers in Biology and Medicine*, 87(February):271–284, 2017.

[3] Deep learning for healthcare: Review, opportunities and challenges. *Briefings in Bioinformatics*, 19(6):1236–1246, 2017.

[4] U. R. Acharya, H. Fujita, S. L. Oh, Y. Hagiwara, J. H. Tan, and M. Adam. Application of deep convolutional neural network for automated detection of myocardial infarction using ECG signals. *Information Sciences*, 415-416:190–198, 2017.

[5] Association for the Advancement of Medical Instrumentation. Testing and reporting performance results of cardiac rhythm and st segment measurement algorithms, 1998.

[6] P. C. Chang, J. J. Lin, J. C. Hsieh, and J. Weng. Myocardial infarction classification with multi-lead ECG using hidden Markov models and Gaussian mixture models. *Applied Soft Computing Journal*, 12(10):3165–3175, 2012.

[7] P. De Chazal, M. O'Dwyer, and R. B. Reilly. Automatic classification of heartbeats using ECG morphology and heartbeat interval features. *IEEE Transactions on Biomedical Engineering*, 51(7):1196–1206, 2004.

[8] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2016-December:770–778, 2016.

[9] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. 2017.

[10] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017*, 2017-January:2261–2269, 2017.

[11] T. Ince, S. Kiranyaz, and M. Gabbou. A generic and robust system for automated patient-specific classification of ECG signals. *IEEE Transactions on Biomedical Engineering*, 56(5):1415–1426, 2009.

[12] S. Kiranyaz, T. Ince, and M. Gabbouj. Real-Time Patient-Specific ECG Classification by 1-D Convolutional Neural Networks. *IEEE Transactions on Biomedical Engineering*, 63(3):664–675, 2016.

[13] B. U. Köhler, C. Hennig, and R. Orglmeister. The principles of software QRS detection. *IEEE Engineering in Medicine and Biology Magazine*, 21(1):42–57, 2002.

[14] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 2:1097–1105, 2012.

[15] A. Lyon, A. Mincholé, J. P. Martínez, P. Laguna, and B. Rodriguez. Computational techniques for ECG analysis and interpretation in light of their contribution to medical advances. *Journal of the Royal Society Interface*, 15(138), 2018.

[16] Y. Ma, X. Ding, Q. She, Z. Luo, T. Potter, and Y. Zhang. Classification of Motor Imagery ECG Signals with Support Vector Machines and Particle Swarm Optimization. *Computational and Mathematical Methods in Medicine*, 2016(5):667–677, 2016.

[17] Y. Ma, X. Ding, Q. She, Z. Luo, T. Potter, and Y. Zhang. Classification of Motor Imagery EEG Signals with Support Vector Machines and Particle Swarm Optimization. *Computational and Mathematical Methods in Medicine*, 2016(5):667–677, 2016.

[18] Massachusetts Institute of Technology. MIT-BIH arrythmia database. www.physionet.org/content/mitdb/1.0.0. Accessed: 2020-06-21.

[19] G. B. Moody and R. G. Mark. The impact of the MIT-BIH arrhythmia database. *IEEE Engineering in Medicine and Biology Magazine*, 20(3):45–50, 2001.

[20] M. M. Rahhal, Y. Bazi, H. Alhichri, N. Alajlan, F. Melgani, and R. R. Yager. Deep learning approach for active classification of electrocardiogram signals. *Information Sciences*, 345:340–354, 2016.

[21] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.

[22] World Health Organisation. The top 10 causes of death. https://www.who.int/news-room/fact-sheets/detail/the-top-10-causes-of-death. Accessed: 2020-06-21.

[23] Y. Xia, H. Zhang, L. Xu, Z. Gao, H. Zhang, H. Liu, and S. Li. An Automatic Cardiac Arrhythmia Classification System with Wearable Electrocardiogram. *IEEE Access*, 6:16529–16538, 2018.

[24] Y. Zheng, Q. Liu, E. Chen, Y. Ge, and J. L. Zhao. Time series classification using multi-channels deep convolutional neural networks. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 8485 LNCS:298–310, 2014.

[25] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le. Learning Transferable Architectures for Scalable Image Recognition. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 8697–8710, 2018.