Detection of Abnormalities in Electrocardiogram (ECG) using Deep Learning

João Pestana[a], David Belo[b] and Hugo Gamboa[c]
LIBPHYS-UNL / FCT, New University of Lisbon, Portugal

Keywords: Electrocardiogram, Signal Processing, Deep Learning, Artificial Intelligence, Arrhythmia Detection, Noise Detection.

Abstract: The Electrocardiogram (ECG) cyclic behaviour gives insights on a subject's emotional, behavioral and cardiovascular state, but often presents abnormal events. The noise made during the acquisition, and presence of symptomatic patterns are examples of anomalies. The proposed Deep Learning framework learns the normal ECG cycles and detects its deviation when the morphology changes. This technology is tested in two different settings having an autoencoder as base for learning features: detection of three different types of noise, and detection of six arrhythmia events. Two Convolutional Neural Network (CNN) algorithms were developed for noise detection achieving accuracies of 98.18% for a binary-class model and 70.74% for a multi-class model. The development of the arrhythmia detection algorithm also included a Gated Recurrent Unit (GRU) for grasping time-dependencies reaching an accuracy of 56.85% and an average sensitivity of 61.13%. The process of learning the abstraction of a ECG signal, currently sacrifices the accuracy for higher generalization, better discriminating the presence of abnormal events in ECG than detecting different types of events. Further improvement could represent a major contribution in symptomatic screening, active learning of unseen events and the study of pathologies to support physicians in the future.

1 INTRODUCTION

In the context of medicine and healthcare in general, physiological signals offer information about the health state. For the ECG, in particular, the morphological and spectral components of each cycle provides hints of the emotional, behavioral and the health state of the individual (Silipo and Marchesi, 1998; Brown, 1999).

With a deeper understanding of pathologies and the development of diagnostic methods and therapies allied to the evolution of technology in the medical field (e.g. wearables), greater healthcare expectations emerged in terms of efficiency. The merged fields of machine learning and the medical field provide technologies that are useful in assisting medical practitioners not only by decision making processes, diagnosis and treatment, but also by continuous monitoring (Johnson et al., 2018; Coiera, 2003; Faust et al., 2018). On this account Deep Neural Networks (DNN) present themselves as a tool that learns higher abstractions of data, while dealing with a large amount of data and the decrease the need for feature engineering (Johnson et al., 2018; Faust et al., 2018).

We propose a framework that learns the abstraction of the default morphology of a normal ECG and detects the divergence when it is modified, in two different scenarios: contamination due to noise and artifacts, and; the presence of symptomatic occurrences.

2 ELECTROCARDIOGRAM

The ECG measures the electrical activity generated by the heart activity in relation to time by inserting electrodes on the skin. This signal is used to diagnose the cardiac health state and can be accessed by understanding the fundamentals of its cycle, i.e. the morphological sequence its characteristic waves: P, QRS complex, T and U (Acharya et al., 2007).

The P wave is due to the depolarization of the atrial myocardium, the following QRS-complex, a fast spiking wave that stimulates the ventricular contraction. The end of the cycle is made by the T wave

[a] https://orcid.org/0000-0002-1760-5255
[b] https://orcid.org/0000-0002-5337-0430
[c] https://orcid.org/0000-0002-4022-7424

Pestana, J., Belo, D. and Gamboa, H.
Detection of Abnormalities in Electrocardiogram (ECG) using Deep Learning.
DOI: 10.5220/0008967302360243
In *Proceedings of the 13th International Joint Conference on Biomedical Engineering Systems and Technologies (BIOSTEC 2020) - Volume 4: BIOSIGNALS*, pages 236-243
ISBN: 978-989-758-398-8; ISSN: 2184-4305

that indicates preparation to a new cycle with the re-polarization of the ventricular myocardium. The U may indicate the late depolarisation of the ventricular myocardium (Acharya et al., 2007). The Signal-to-Noise Ratio (SNR) comprises a heart rate between 60 and 100 bpm at rest while maintaining this sequence.

2.1 ECG Noise

Noise can be defined as any signal that is not related to the heart's electrical activity that obstructs the ECG signal through interference (Rodrigues et al., 2017). Usually, the signal is obtained in a controlled environment, although the usage of ambulatory ECG is more associated with certain variables that influence the noise (Xiong et al., 2019).

Noise interference has many possible sources, which can be classified as biological, like motion artifact and muscle contraction, and environmental or non-physiological noise, related to nearby electrical devices and variables associated with the equipment (Rodrigues et al., 2017).

baseline-wander (BW) artifacts are caused by movement such as the contraction and relaxation of the thoracic cage muscles, which is evident if the event is periodic. Noise affected signals can also originate from the misplacement of electrodes and loss of contact between the sensor and the skin, called electrode motion (EM). Additionally, the activation of the muscles cause the EMG signal interference in the ECG signal, revealing motion artifact (MA).

2.2 Arrhythmias

Arrhythmias can be defined as rhythms that are considered abnormal, considering the Normal Sinus Rhythm (NSR), that are related to malfunctions in the heart's electrical activity and changes in the heart tissue (arr,). They can be detected in ECG by analysing the heart rate, the presence or absence of certain waves, the duration of intervals, wave overlap, abnormal timing of cardiac events, the shape, size and direction of the waves, between others.

Common physiological arrythmias can be detected from variations of the intervals between R peaks (RR intervals). These changes are correlated with the respiration rhythm and the sympathetic nervous system. When RR interval is lower than 60 bpm the arrhythmia is named tachycardia and higher than 100bpm bradycardia. (arr,).

The pathophysiological arrhythmias may be caused due to blocks in the electrical impulse conduction, enlargement of the myocardium, pericarditis, electrolyte imbalance, respiratory diseases, drugs, hypothermia and accessory conduction pathways(Acharya et al., 2007).

3 STATE OF THE ART

Over the years, various methods have been studied for ECG noise and arrhythmia detection. The following sections show an overview of the general developments in these areas.

3.1 Noise Detection in ECG

A decision rule-based algorithm is proposed by (Satija et al., 2018) that calculates the maximum absolute amplitude, number of zero crossings and local maximum peak amplitude of the autocorrelation function are calculated, after using a modified ensemble empirical mode decomposition. This work was able to discriminate between six signal groups with an accuracy of 98.93%. (Ansari et al., 2018) proposed a 16 layer Convolution Neural Networks (CNN) which predicts once per second on a 10 second inputs. An AUC of 0.977 was reached for this binary classification model with a 88,7% sensitivity (John et al., 2018). Ansari et al. (2018) also developed a CNN to detect usable and unasable ECG segments, in terms of calculating the heart rate variability (HRV). The filtering using the CNN resulted in an area under curve (AUC) of 0.96 for the classification of noise affected segments, compared to an AUC of 0.87 for a Support Vector Machine model.

3.2 Pathological Event Detection in ECG

More recently, (Acharya et al., 2017) used a CNN composed of 11 layers that analyse with 2 or 5 seconds for detection of atrial fibrillation, atrial flutter, and ventricular fibrillation. This study achieved accuracy, sensitivity and specificity values of 92.50%, 98,09% and 93,13%, respectively, were achieved, and for the two seconds ECG segments and 94,90%, 99.13% and 81.44% in the same order.

Also, a combination of CNN with Recurrent Neural Networks (RNN) was developed by Andersen et al. (2018) in order to distinguish between atrial fibrillation and normal segments for ECG recordings of 24 hours. The extracted features using the CNN module were processed by the RNN achieving a sensitivity of 98.98% and a specificity of 96.95%, making predictions for 24 hours of ECG signal in under one second. This algorithm was sensible to noise, reducing significantly its results when present (S. Andersen et al., 2018).

Recently, Hannun et al. (2019) developed a 34 layer DNN that could detect between 10 different arrhythmias and NSR using raw signal, while also detecting noise corrupted segments, constituting 12 prediction classes in total. The results were higher than average cardiologists in terms of sensitivity and matching specificity values. The average AUC was of 0.97 and the confusion matrices were similar, accentuating the same problematic rhythm classes for both (Hannun et al., 2019).

4 METHODS

Before the signals are fed to the deep learning algorithms, preprocessing steps are implemented in order to transform the raw data. These are imperative for the models to learn the underlying patterns, to provide better generalization and overall quality. Therefore, the signal is submitted to the decimation method in order to reduce the number of samples, while mitigating the loss of information. This technique applies a low pass anti-aliasing filter (order 8 Chebyshev type I) that suppresses all frequencies that may cause aliasing, before the subsampling process.

In order to reduce the baseline noise that is contained in the normal ECG, the clean signal was convolved with a Hanning window. All records were normalized and the average was removed, in order to be interpretable by the network, with the following rule:

$$x' = \frac{x - \bar{x}}{max(x) - min(x)} \tag{1}$$

where the normalized signal is x' and x denotes the raw signal. Finally the signals are segmented in windows of 64 samples, since the mean duration of a cycle is of 1 second, most of the widows have one QRS complex.

4.1 Architectures

In order to perform ECG classification according to the two tasks at hand, two different architectures were chosen and optimized to classify noise affected segments and to classify different types of arrhythmia and NSR. Since autoencoders are good in learning general features from the input, they were selected as the base for the implementation of both detection algorithms.

In a first stage, an autoencoder will lean the characteristics of a normal ECG signal in the purest form possible, i.e. with a good SNR and no symptomatic events. Consequently, the latent vector will contain the coded information of the input, and thus also called ”feature vector”. When this structure is fed with a modified ECG it will produce a different latent space. As the state changes it can be analysed and detected using a classifier.

The training of each algorithm was performed using batches of 256 shuffled windows, each without an overlap for the noise detection experiment, and 50 % for the arrhythmia. The testing was made with the rest of each signal. Since the normal ECG cycle occurs during one second it is expected that each window contains at least one QRS complex.

4.1.1 Autoencoder

CNN autoencoders are able to encode the input data, as well as reconstruct it at the output layer with the goal of reproducing the input signal as similar as possible. As the convolutional layers are capable learning feature maps, which extracts the essential features to reconstruct the input, these algorithms are an unsupervised feature learning mechanism (Mao et al., 2017).

The diagram for the encoder is depicted in Fig. 1, which comprises three convolutional blocks, each containing one convectional layer with Rectified Linear Unit (ReLU) activation and a max pool, producing the final latent vector with 256 elements.

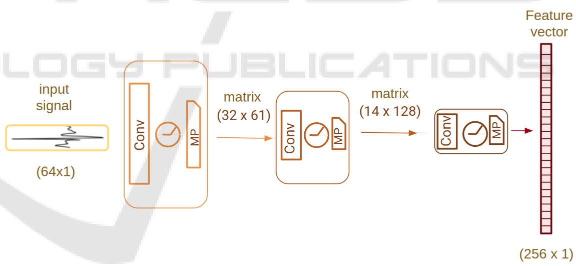

Figure 1: The encoder is composed by three blocks composed by a convolutional (”Conv”) with a ReLU activation and a max pooling ”MP” layers.

Even though the decoder won’t be used specifically for the classification of abnormal events in ECG, this component is vital for the training phase. The training, monitoring and optimization of the encoder is made by training comparing the input and output with the mean squared error and RMSProp as the optimizer (initial learning rate of 0.001).

4.1.2 Noise Detection Neural Network

Two different detection networks were developed: (1) a binary noise detection model capable of classifying between Noise Affected Signals (NAS) or Normal Signal (NS); (2) multi-class detection using three types commonly found in ECG: BW, MA and EM noise.

Both algorithms follow the same structure, as depicted in Fig. 2. The difference between the algorithms is in the number of layers in the fully connected (FC) network because of the number of classes. As this network progresses toward the end, each layer decreases by half the number of neurons, starting in 256, with ReLU activations in the first layers and softmax in the last one, and since the last layer contains the number of classes for each system, the first settled with 9 while the second with 10 layers. In both setups, a dropout of 50% is implemented in the fifth layer. The training was also made with RMSProp, but with cross entropy loss function.

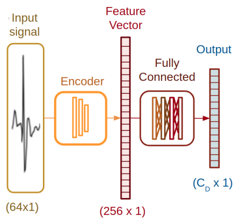

Figure 2: Model architecture for the noise detection, with C_D number of classes.

4.1.3 Arrhythmia Detection Neural Network

The sequential approach for detection and prevention of the arrythmia events is portrayed in the diagram in Fig. 3. It is composed by an an encoder, FC layers and a RNN comprising three sequential Gated Recurrent Units (GRU). A RNN is a particular type of neural network that can be described as a dynamic sequential data processing model where its current internal state is affected by the previous one (Pascanu et al., 2013). The GRU is a particular type that manages the learning process using gates (Ghimes et al., 2018).

After the ECG features from each widow are extracted using the previous CNN encoder, N number of feature vectors are stacked, resulting in matrices of size ($N \times F_D$), where F_D is the dimension of the feature vector. Since the RNN module has its own set of hidden units (H_D), the FC layer, with linear activation, makes the connection between both modules. The parameters used for this experiment were $F_D = 256$ and $H_D = 128$ and $N = 32$

After the three GRU layers the output is carried into a FC layer, classifying into one of the seven classes ($C_D = 7$). The final module starts with a batch normalization layer but is similar to the previous FC layers used for classification of the noise algorithms from $H_D = 128$, until the size of $C_D = 7$.

The training of this algorithm was made with RMSProp with the initial learning rate of 0.001, which was automatically adjusted during its training phase. The input data is separated into batches of 256 signal windows ensuring that it was within the computational limitations.

4.2 Datasets

All the datasets used are public and can be accessed through Physionet, a biosignal database that was created under the auspices of the National Institutes of Health (Goldberger et al., 2000).

The Fantasia dataset was acquired from twenty young (21-34 yr) and twenty elderly (68-81 yr) while exposed to 120 min of continuous supine resting ECG recordings (250 Hz sampling rate) while watching the Disney's movie Fantasia (Iyengar et al., 1996). The signals collected by this database are mostly clean and free from symptoms, which reflects a good baseline for training the autoencoder.

For the noise detection model, both Fantasia and MIT-BIH Noise Stress Database were considered (Moody et al., 1984). The last includes 12 half-hour ECG recordings and 3 half-hour recordings typical noise in ambulatory ECG recordings with 250Hz.

For the last task, the MIT-BIH Arrhythmia Database was used, displaying a sampling frequency of 360 Hz. The chosen 6, out of 14, had enough number of episodes to provide a good balance between classes while training the algorithm: Atrial Fibrillation (AFIB), Atrial Flutter (AFL), Ventricular Bigeminy (B), Paced Rhythm (P), Wolff-Parkinson-White Syndrome, or Pre-excitation (PREX), and Sinus Bradycardia (SBR) (Moody and Mark, 2001).

5 RESULTS

All models were developed, trained and tested using Keras API on top of the Tensorflow library. The hardware included a NVIDIA GeForce GTX 960.

5.1 Autoencoder

The autoencoder was trained and tested using all 40 individuals in the Fantasia Database, being that 70% were used for training and 30% for testing. After approximately 2100 epochs of training the minimum reached loss was 0.0001. For the testing phase, the mean squared error was calculated for each signal reaching a mean value of 0.0026 and a standard deviation of 0.0012 during ±140 minutes. The results for

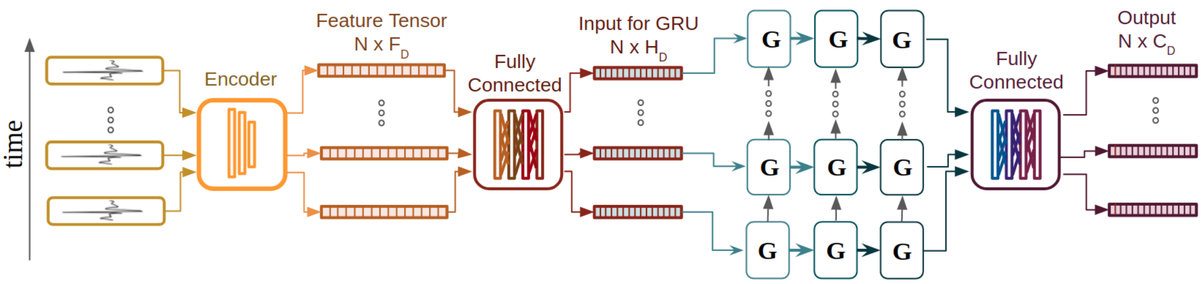

Figure 3: Model architecture for the arrhythmia detection, with $C_D = 7$ classes.

the first 10 individuals of the dataset, which have ages between 21 and 34 years old, are depicted in Fig. 4.

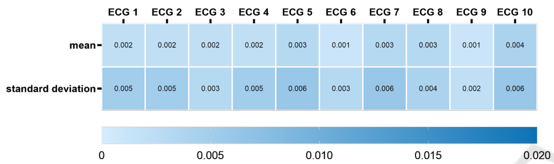

Figure 4: Mean and standard deviation values for the first 10 subjects within the age group of 21 to 34 years old.

The low values exhibited by the results suggest that the model was able reconstruct the main characteristics of the signal, that can be seen in the example shown in Fig. 5.

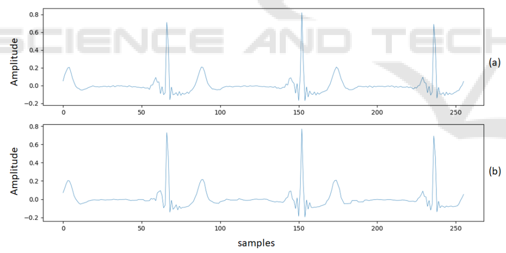

Figure 5: Portion of the signal ECG 9, from the Fantasia Database (a) and reconstruction of the same signal (b).

After analyzing the results, it was concluded that the encoder could be applied for the following classification models.

5.2 Noise Detection Neural Network

The training set for both noise detection models the data was separated in 70% for training and 30% for testing. Furthermore, the testing set had three individuals that the training set did not contain. Both algorithms took ±120min to train.

5.2.1 Binary Noise Detection Model

As stated before, the binary noise detection model classifies ECG segments according to two classes: NS or NAS. The training proceeded until 600 epochs, where an abrupt increase of error was observed, reaching a minimum cross entropy of 0.51 with an accuracy of 98,56%.

The classification results, shown in Table 1, provides evidence that the model was able to successfully detect the presence of noise in the ECG. Further support for this claim is presented in the confusion matrix exhibited in Fig. 6.

Table 1: Binary noise detection model: classification performance (%).

Accuracy	Sensitivity	Specificity
98,18	98,21	98,15

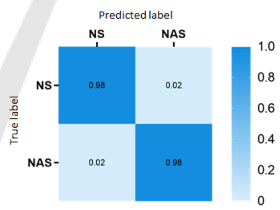

Figure 6: Normalized confusion matrix of the dataset used for the training phase, where NS is the positive label and NAS is the negative label.

The encoded data for both classes were submitted to t-Distributed Stochastic Neighbour Embedding (t-SNE) in order to visualize the proximity between the feature vectors. Therefore, each point in the resulting graph represents the tensor created by the encoder originated by a single window input. Both classes are clearly separated in two clusters when each time window is (equally distributed by label) submitted to t-SNE(Fig. 7). This confirms that the encoder extracts different values for each feature represented by the two classes.

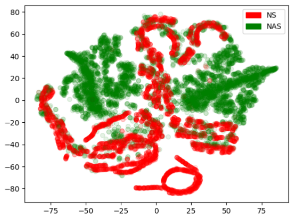

Figure 7: Binary noise detection algorithm - t-SNE representation of normal signal (NS) (red) and Noise affected signal (NAS) (green) encoded data windows.

5.2.2 Multi-class Noise Detection Model

During the training phase, a minimum cross entropy error of 0.85 was reached after 600 epochs, resulting in a training accuracy of 86,17% after ±180min. However, testing accuracy was only able to reach 70,74%.

The evaluation of this model for each class are conveyed in Table 2.

Table 2: Multi-class noise detection model: classification performance for each class (%)

Class	Sensitivity	Specificity
NS	89,77	96,63
BW	72,70	84,26
EM	65,19	91,35
MA	59,11	89,47

By inspecting the normalized confusion matrix in Fig. 8, it is can be stated that even though the algorithm is able to correctly identify normal signal most of the times, it is harder to correctly separate the identity of each type of noise. However, the model was able to discern between EM and MA characteristics.

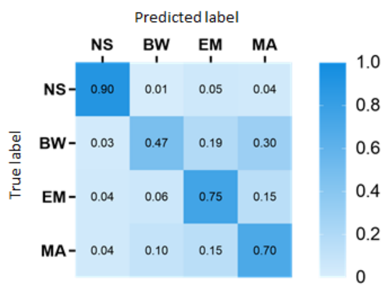

Figure 8: Multi-class Noise detection algorithm - Confusion matrix for the classes: Normal Signal (NS), baseline-wander (BW), electrode motion (EM) and motion artifact (MA).

Comparing these results with the previous experiments, we come to the conclusion that even though it is possible to detect noise corrupted ECG with high accuracy, but differentiating between the several types of noise reveals to be a harder task. This may be due the fact that most of the times signals are affected by different types of noise sources at once and because the encoder was not trained to deal with these differences.

5.3 Arrhythmia Detection Model

Cross-validation was performed with 90% for training the data and the rest for testing. The minimum cross entropy error of 0.65 was reached after 1750 epochs, ending this training phase with a 90,74% accuracy. However, for the test data, the accuracy was only of 56,85%.

The confusion matrix, displayed in Fig. 9, shows that the model was able to correlate between certain characteristics with the correspondent arrhythmia, but underperforming between distinguishing some of them. This DNN architecture had a high classification rate for the P in contrast with the detection of AFIB, often confusing its morphology with AFL, since both of them are quite similar.

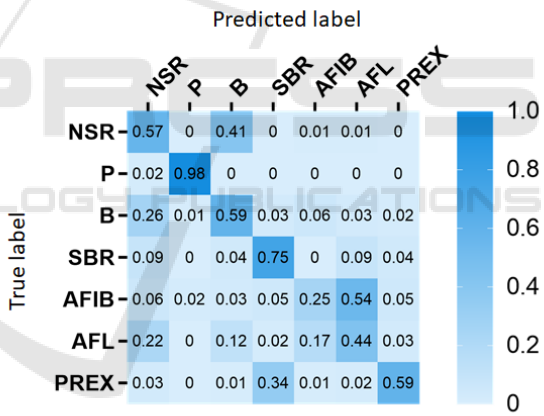

Figure 9: Arrithmia detection algorithm - Confusion matrix for the classes of Atrial Fibrillation (AFIB), Atrial Flutter (AFL), Ventricular Bigeminy (B), Paced Rhythm (P), Pre-excitation (PREX), and Sinus Bradycardia (SBR).

These results are promising and reveal enough facts to embrace the development of this algorithm by fine-tuning its parameters and have some considerations in the training phase of the autoencoder. Unfortunately at this current stage it cannot be used in a real life context, since the model often considers pathological ECG segments as NSR, giving too may false positives and negatives. By analysing the classification performance of each class presented in Table 3, it is possible to conclude that the algorithm often struggles to correctly identify ECG segments with B, AFL and AFIB, reaching an average sensitivity and specificity for this model of 61.13% and 93.10%, respectively. But we are confident that these results can be improved with more powerful tools. Options for these improvement are the increase of the window size, that would help the algorithm understand the differences of AFL and AFIB that, between other aspects, change the distance between peaks and extending the period of time, including more cycles, would make these easier to perceive. In a future work, this algorithm will be tested with a more powerful hardware setting to be able to encompass this parameter change.

Table 3: Arrhythmia detection algorithm: classification performance for each class (%).

Class	Sensitivity	Specificity
NSR	63,35	85,49
P	97,12	99,73
B	36,31	93,67
SBR	62,60	96,33
PREX	81,33	94,45
AFL	38,14	91,88
AFIB	49,06	90,12
AFIB+AFL	84,80	90,91

Hannun et al (2019) were able to reach an average sensitivity of 75.22% for a 12 class model, compared to the average sensitivity of 61.13% for the proposed model. However, is it important to note that in the mentioned study, AFIB and AFL classes were combined into one unique class (Hannun et al., 2019). Figure 10 represents the merged results for a direct comparative study Figure 9, and when the sensitivity and specificity values are re-calculated (Table 3) the first mean increases to 84,80%, but the second average decreases 70,91%, a value that only differs from the study mentioned by 4,31%.

6 CONCLUSIONS

As for the last example of application of DNN architectures, the objective was to develop a framework that detects abnormalities in the normal pattern of the ECG signal, by creating a model that learns the key characteristics of a normal ECG cycle. The resultant models were capable of detecting two types of abnormalities: noise and pathological events.

A machine-learned extractor was developed in order to be included in the architectures for noise and arrhythmia detection, by using the encoder module of a trained autoencoder. From the shown results it is suggested that by successfully reproduce the input signals, the autoencoder learns the key characteristics of a normal ECG cycle.

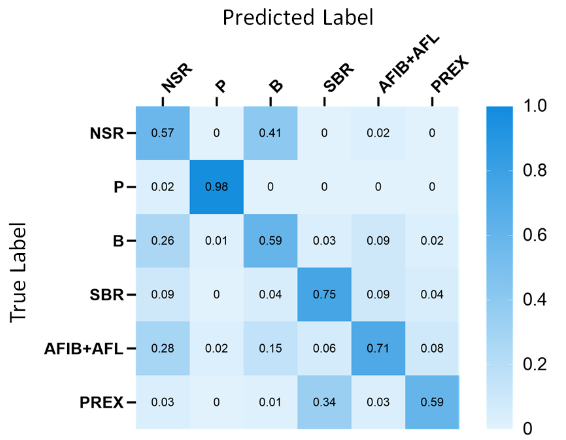

Figure 10: Arrhythmia detection algorithm -Normalized confusion matrix of the dataset used for the testing phase, with classes: Atrial Fibrillation and Atrial Flutter (AFIB + AFL), Ventricular Bigeminy (B), Paced Rhythm (P), Pre-excitation (PREX), and Sinus Bradycardia (SBR).

For the noise detection models, reaching state-of-the-art performance in the case of binary classification, despite achieving lower results in discriminating different types of noise. In the case of the arrhythmia decision system, the model did not perform as well as the compared state-of-the-art accuracy, but the sensitivity and specificity showed to be promising.

The improvement of the system could rely on the need to increase the power for the encoder by supplying it with more ECGs. As the autoencoder learns and replicates an increasing quantity of different segments of asymptomatic signals, the more this algorithm will embrace the abstract notion of what is an ECG. The presence of more noisy and symptomatic episodes would also contribute to the classification blocks, since the dataset was limited in the number of provided segments.

In sum, the performance of these algorithms proves that it is possible to detect generic abnormal events in ECG using the proposed architectures. Therefore, the implementation of an encoder for classification algorithms proves to be effective in the detection of abnormal events by learning the signals key characteristics.

As for the detecting each type of abnormal events the model did not reach state-of-the-art results as expected since the algorithms developed sacrifice accuracy to reach better generalizations for all ECG signals to be applicable in real life. Applications that understand the deviation from the normal signal into unseen pathologies and the possibility of creating new classes that did not exist before, could be extremely beneficial for future not only for ECG, but also other physiological signals.

One example is the application of this framework to active learning mechanisms as most of the available algorithms are specific to specific conditions, while not detecting others. The ability to detect similar abnormalities in the ECG signal, the classifier could actively create new labels for patterns that diverge from the normal ECG signal, while increasing new diagnostic possibilities. That said, this technique could not only aid in the diagnostic of arrhythmias, but also contribute to the study of these pathological anomalies, by finding correlations between the different expressions of the ECG morphology, and also by exploring these manifestations with different variables in mind such as the patient's gender, age, medications, among other factors. Another example is by analysing billions of data points in wearable data, the medical doctor could focus only in the parts which contained different aspects of the signal, so that he could diagnose without wasting hours or days of analysis.

This method helps to increase the possibilities, not only for detection, but also for studying what exactly is a normal cycle and which are the deviation patterns. By having early detection procedures, one could early seek for medical help without developing a symptomatic episode. This preventive point of view could represent a major change in perspective in how medical care should be delivered in the future.

REFERENCES

Arrhythmia — National Heart, Lung, and Blood Institute (NHLBI).

Acharya, R., Krishnan, S. M., Spaan, J. A., and Suri, J. S. (2007). *Advances in cardiac signal processing*. Springer.

Acharya, U. R., Fujita, H., Lih, O. S., Hagiwara, Y., Tan, J. H., and Adam, M. (2017). Automated detection of arrhythmias using different intervals of tachycardia ECG segments with convolutional neural network. *Information Sciences*, 405:81–90.

Ansari, S., Gryak, J., and Najarian, K. (2018). Noise detection in electrocardiography signal for robust heart rate variability analysis: A deep learning approach. In *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 5632–5635. IEEE.

Brown, B. H. (cop. 1999). *Medical physics and biomedical engineering*. Medical Science Series. Institute of Physics Publishing, Bristol.

Coiera, E. (2003). *Guide to Health Informatics*. Oxford University Press, second edition.

Faust, O., Hagiwara, Y., Hong, T. J., Lih, O. S., and Acharya, U. R. (2018). Deep learning for healthcare applications based on physiological signals: A review. *Computer Methods and Programs in Biomedicine*, 161:1 – 13.

Ghimes, A.-M., Avram, A.-M., and Vladuta, V.-A. (2018). A character prediction approach in a security context using a recurrent neural network. In *2018 International Symposium on Electronics and Telecommunications (ISETC)*, pages 1–4. IEEE.

Goldberger, A. L., Amaral, L. A., Glass, L., Hausdorff, J. M., Ivanov, P. C., Mark, R. G., Mietus, J. E., Moody, G. B., Peng, C.-K., and Stanley, H. E. (2000). Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals. *Circulation*, 101(23):e215–e220.

Hannun, A. Y., Rajpurkar, P., Haghpanahi, M., Tison, G. H., Bourn, C., Turakhia, M. P., and Ng, A. Y. (2019). Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. *Nature Medicine*, 25(1):65–69.

Iyengar, N., Peng, C., Morin, R., Goldberger, A. L., and Lipsitz, L. A. (1996). Age-related alterations in the fractal scaling of cardiac interbeat interval dynamics. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 271(4):R1078–R1084.

John, J. N., Galloway, C., and Valys, A. (2018). Deep convolutional neural networks for noise detection in ecgs. *arXiv preprint arXiv:1810.04122*.

Johnson, K. W., Soto, J. T., Glicksberg, B. S., Shameer, K., Miotto, R., Ali, M., Ashley, E., and Dudley, J. T. (2018). Artificial intelligence in cardiology. *Journal of the American College of Cardiology*, 71(23):2668 – 2679.

Mao, Z., Yao, W. X., and Huang, Y. (2017). Eeg-based biometric identification with deep learning. In *2017 8th International IEEE/EMBS Conference on Neural Engineering (NER)*, pages 609–612.

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

Moody, G. B., Muldrow, W., and Mark, R. G. (1984). A noise stress test for arrhythmia detectors. *Computers in cardiology*, 11(3):381–384.

Pascanu, R., Gulcehre, C., Cho, K., and Bengio, Y. (2013). How to construct deep recurrent neural networks.

Rodrigues, J., Belo, D., and Gamboa, H. (2017). Noise detection on ecg based on agglomerative clustering of morphological features. *Computers in Biology and Medicine*, 87:322 – 334.

S. Andersen, R., Peimankar, R., and Puthusserypady, S. (2018). A deep learning approach for real-time detection of atrial fibrillation. *Expert Systems with Applications*, 115.

Satija, U., Ramkumar, B., and Manikandan, M. S. (2018). Automated ecg noise detection and classification system for unsupervised healthcare monitoring. *IEEE Journal of Biomedical and Health Informatics*, 22(3):722–732.

Silipo, R. and Marchesi, C. (1998). Artificial neural networks for automatic ecg analysis. *IEEE Transactions on Signal Processing*, 46(5):1417–1425.

Xiong, F., Chen, D., Chen, Z., and Dai, S. (2019). Cancellation of motion artifacts in ambulatory ECG signals using TD-LMS adaptive filtering techniques. *Journal of Visual Communication and Image Representation*, 58:606–618.