Transfer Learning for tuberculosis screening by single-channel ECG

Valeriia Guryanova
Faculty of Computational Mathematics and Cybernetics
Lomonosov Moscow State University
Moscow, Russia
kibitovavaleria@gmail.com
ORCID: 0000-0002-3613-8103

***Abstract*—Tuberculosis is one of the leading causes of death in the world. The majority of the population is not able to regularly conduct specific examinations, such as x-ray examinations, for the presence of tuberculosis. Currently, there are mobile devices for measuring ECG, which allow taking measurements without leaving home. This article explores the possibility of determining tuberculosis based on a single-channel mobile ECG. One of the general top-performance neural networks is used as a classifier. This article also explored the possibility of such a classification based not on raw data, but the generated image. The image interprets the prediction of the neural network and makes it possible for the doctor to understand the model's decision better. The article shows the promising performance and provides proof of concept of such screening. Different ratios of precision and recall are provided, which can be adjusted depending on the situation.**

***Keywords*—Neural Networks, ECG classification, Tuberculosis, Deep Learning, Transfer Learning**

I. Introduction

Tuberculosis (TB) is caused by bacteria (Mycobacterium tuberculosis) that most often affect the lungs. Tuberculosis is curable and preventable. TB is easily spread from person to person through the air. TB is one of the ten leading causes of death in the world [1]. It is crucial to determine tuberculosis in time to prevent its spread and begin treatment of an ill person. It takes a long time for both the physician and the patient to carry out specific tests for TB (e.g., such as X-rays or Xpert MTB/RIF), and as such, they cannot be conducted with relatively high regularity. So, it would be useful to have some pre-screening, which could be done quite often and which would allow identifying people with a high chance of having tuberculosis and conduct specific tests only for them.

ECG is a signal that displays the electronic activity of the heart. Each ECG recording shows the potential difference between two electrodes located on the surface of the body. Each of the measured potential differences is called a lead. In medical institutions, 12-lead ECGs are commonly used. Currently, on the market, there are mobile devices that are capable of reading an electrocardiogram (ECG) of a person, for example, AliveCor [2], CardioQvark [3]. Such devices read only one of the 12 leads. Such devices can be used as often as required. Anybody can have just a few devices per institution or organization to conduct such pre-screenings every day. People who have a high chance of having a disease will need to see a TB specialist as soon as possible.

The motivation for this study was the work that shows the relationship between cardiological and tuberculosis diseases [4]. Besides, often the tuberculosis bacteria themselves affect the heart and thus affect its electrical activity. At the moment, this is the first work that is devoted to the problem of tuberculosis detection via single-channel mobile ECG.

Now neural networks are one of the most popular methods for analyzing medical signals and images. There are plenty of works in which medical images and ECGs are analyzed using neural networks (NN), for example, [5], [6], [7], [8], [9]. Usually, raw ECGs are analyzed using recurrent [10] or convolutional neural [8] networks with 1d convolutions. In this paper, it is proposed to analyze an ECG visualization and not a raw signal. So the image but not a signal is analyzed. There are some works in which neural networks analyze the image of an ECG signal [11] [12] and show good performance. However, in these articles, architectures were trained from scratch and were specially selected for the task. There are large number of pretrained architectures for images [13], [14], [15]. Since, for many tasks, pretrained architectures show significantly better performance than architectures trained from scratch [16], there is an assumption that they will show better performance on the paper's task. There is work [17] that shows that pretrained architectures can be used as a feature extractor for phonocardiogram signals (PCG), which also provides additional motivation for exploring the possibility of using pretrained architectures.

This paper aims to investigate the possibility of determining tuberculosis from ECG images, the possibility of using pretrained models for such tasks, and of the interpretation of such models. The main scientific novelty of this paper is the use of pre-trained architecture for ECG classification. It is also the first article that explores the feasibility of determining tuberculosis by a single-channel ECG and explores the relationships of metrics that a solution can achieve.

The paper is organized as follows. Initially, the dataset and data preprocessing methods are described. Then, the process of generation of the ECG image and the architecture of the neural network are given. Then the results of experiments and interpretation of the neural network are reported. Finally, the limitations of the article are analyzed, and the conclusion is given.

II. Dataset

All data used in this study was collected using the Cardio-Qvark device [3]. The dataset consists of 1232 ECGs from 136 people with TB and 3609 ECGs from people without TB. Dataset sex distribution among people with tuberculosis and people without tuberculosis is shown in Fig.1. Dataset age distribution among people with tuberculosis and people without tuberculosis is shown in Fig.2.

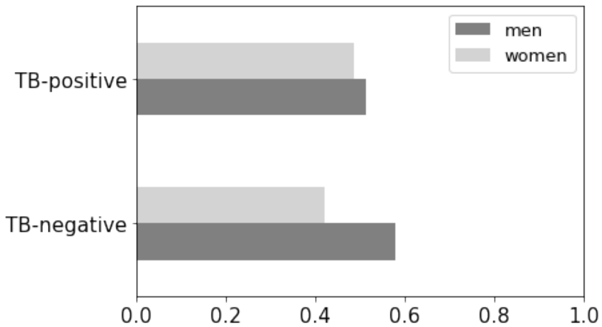

Fig. 1. Dataset sex distribution.

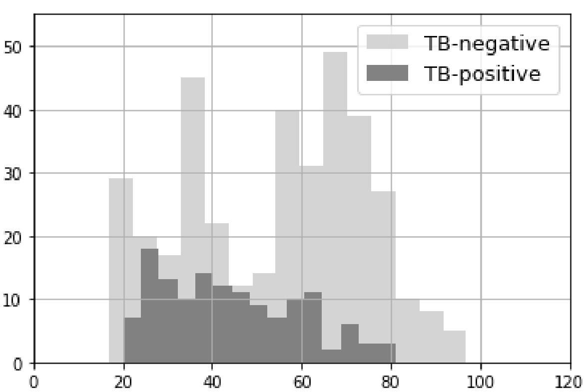

Fig. 2. Dataset age distribution.

ECG recordings were sampled as 1000 Hz, and they have been filtered by the CardioQvark device. The length of ECG recordings varies from 30.5 seconds to 900 seconds. For trend subtraction, the median filter with a kernel size of 187 was used.

III. Image Generation Process and Neural Network Architecture

Signals were resampled to a frequency of 200 Hz via the polyphase resample method from the SciPy Python library [18] to reduce the image size. Signals longer than 100 seconds were divided into sub-signals. The signal is divided sequentially into sections of 100 seconds to receive the sub-signals, if the length of the last section is less than 30 seconds, then it was not added to the sample set.

Since the signals were taken from a device that can have noise, and the signal could disappear for various reasons, it is necessary to remove signals that contain much noise or do not contain any useful information. R-peaks that represent heartbeat were calculated using the Pan–Tompkins algorithm [19]. Then the distances between the two peaks were calculated. It was measured heuristically that the signal section almost certainly does not contain useful information if the distance between the two R-peaks is more than 2.6s or less than 0.3s. The signal was removed from the dataset if there were more than 55 percent of such distances in it.

After dividing the long signals into sub-signals and removing the signals that contain mainly noise, 13776 signals were obtained. Ten thousand ninety-seven of them belong to 367 TB-positive people, and 3679 of them belong to 136 TB-negative people. More than 90 percent of the resulting signals is 100 seconds long. The number of signals per TB-positive patients varies from 1 to 1540. The number of signals per TB negative patients varies from 3 to 198.

The image was generated using the Matplotlib plot function [20]. The whole ECG signal was divided into seven parts of 2857 points each, with dots per inch of 70. The size of the figure was set to 7.4 by 7.4 inches. Since the matplotlib generates an image in the RGBA format, and most neural networks use RGB format, only the 4th channel was used. This channel was multiplied to create three channels. The example of the generated image is given in Fig. 3.

Fig. 3. Generated ECG image.

ResNext 101 [15] was used as the main model for the experiments. The main idea of this model is the introduction of the new dimension, which authors called "cardinality" - the size of the set of transformations. The main function of this dimension is the control over complex transformations. Even while maintaining complexity, increasing ”cardinality” can improve classification accuracy. In this paper, a model with a ”cardinality” value of 32 was used. This model is one of the top best models for classification on ImageNet [21] and is easy to use and retrain on new data.

This model has four blocks of convolutional layers. The following neural network configurations were used to explore the possibilities of transfer learning: the completely untrained network (UN), pretrained on ImageNet network with no frozen layers (PN), pretrained on ImageNet network with frozen layer 1 (or 2, 3, 4) and all layers before it (L1(2, 3, 4)). In the case of freezing up to layer four, only one fully connected classifier layer is trained.

Augmentation is a method in which artificially generated data is added to the original data. This method is one of the methods that is widely used to improve the quality of neural network models [22]. Online augmentation is a method that replaces the part of the data to the synthetic one during the training stage. A study was conducted that investigated online augmentation for the classification quality improvement of single-channel ECGs by using a neural network [23]. The data that was used in that article contains the dataset discussed in this study. However, that study did not use the representation of the signal as an image, nor did it study the possibilities of such training. Research in an article on online augmentation has shown that it helps in improving the quality of classification of single-channel signals.

One of the most promising methods was the ResampleParts method. This method resamples parts of the signal with a specific resampling coefficient using the polyphase resample method from the SciPy Python library [18]. The specific size of the signal part and resampling coefficient is set at random each time from a certain interval. Due to the reason mentioned above, this augmentation method was chosen to increase the quality of tuberculosis detection.

IV. Experiments and Results

All experiments were conducted using Python language. PyTorch [24] was used as the main neural network framework. Signals were read and transformed using NumPy [25] and SciPy [18] libraries.

To correctly evaluate the results, the entire dataset was divided into training, test, and validation set. The training set was used to train the model. The validation set was used to select the hyperparameters and for early stopping. A test set was used to check the final performance. The full dataset was partitioned into three parts in a specific way, such that each of the three groups contained different patients. The train set contains 2648 original ECGs from 247 patients. Among them, 69 patients with 689 ECGs have TB. After the preprocessing stage the train set has 7495 ECGs and among them 2058 from TB-positive persons. The validation set contains original 1062 ECGs from 112 patients. Among them, 30 patients with 255 ECG have TB. After the preprocessing stage, the validation set has 3076 ECGs and among them 765 from TB-positive person. The test set contains original 1130 ECGs from 144 patients. Among them, 37 patients with 288 ECGs have TB. After the preprocessing stage, the validation set has 3205 ECGs and among them 857 from TB-positive persons. The number of patients after the preprocessing stage stays the same in all datasets.

For reproducibility, all neural network configurations were trained with a random state of 10. This random state was fixed for NumPy, PyTorch CPU, and GPU variables and also for standard Python random library. Also, the CUDA state was set to deterministic. For all configurations, Adam was used as an optimizer, and batch size was set to 16. For the untrained network, the learning rate was set to 0.0005 and is dropped by a factor of 0.0001 after the 8th epoch. For the pretrained network, the learning rate was set to 0.0005 and is dropped by a factor of 0.001 after the 8th. For the pretrained network with unfrozen and frozen layers, the learning rate was set to 0.0001 and is dropped by a factor of 0.001 after the first and fourth epoch. All parameters that are not described were set to default values of PyTorch library [24]. The metric that was used for early stopping was the area under the precision-recall (PR AUC) curve since classes are not balanced in samples. The precision-recall curve illustrates the ratio between precision and recall of different thresholds. A high value of precision indicates that a model has a low false-positive rate, and a high value of recall shows that model has a low false-negative rate. A high value of the PR AUC corresponds to both high recall and high precision.

Another important metric that is useful to take into consideration during model evaluation is an area under the ROC curve (ROC AUC). ROC curve shows the ratio between the true positive rate and the false positive rate. This metric could be useful as it shows the probability that the model ranks a random positive example higher than a random negative example. However, this metric can be sensitive to unbalanced classes in data, and so it should not be the only metric for performance evaluation. Table I shows the results from different network configurations via the ROC AUC score and PR AUC score.

TABLE I. Experiment Results

Quality Criteria	**Network Configuration**					
	UN	***PN***	***L1***	***L2***	***L3***	***L4***
ROC AUC	0.9079	0.93	**0.934**	0.9307	0.9138	0.8089
PR AUC	0.7956	0.83	**0.84**	0.8399	0.8081	0.615

As can be seen from table I, the best performance is obtained with configurations L1 and L2. This table shows that transfer learning is useful for classifying ECG signals. These results also correspond to the logic that the first layers of a neural network learn simple dependencies such as lines and simple shapes. When freezing a neural network up to the last layers, a worse performance can be gotten than on a network without pretraining (UN), since the features on the last layers are not relevant for this specific task. Thus, based on metrics, L1 can be considered the best model.

For the improvement of this best model earlier described, online augmentation was used. The interval for the resampling part was set to $(40, 200)$. The two types of resampling were used: frequency decreasing (FD) and frequency increasing (FI). The resulting sample rate is calculated as the upsample factor divided on the downsampling factor. For the frequency, increasing the upsample factor was getting from 2 to 3, and the downsampling factor was gotten from 3 to 5. For the frequency decreasing upsample, factor was gotten from 3 to 6, and the downsampling factor was gotten from 1 to 3. The probabilities of augmentations was chosen from values $[0.1, 0.2, 0.3, 0.4, 0.5]$. The best model was determined based on the PR AUC metric on the validation set. The best model was the model with the probability of augmentation of 0.5. The best model has 0.9475 ROC AUC and 0.875 PR AUC on the test set. Based on the results of the experiments, it is reasonable to conclude that augmentation improves model performance.

In most real-life scenarios, the probability of certain events can not be used to make a decision. A specific label is assigned to each patient using probability thresholding. To explore the balance between precision and recall in a situation of unbalanced classes, F_β score metric is used:

$$F_\beta = (1 + \beta^2)\frac{\text{precision} \times \text{recall}}{\beta^2\text{precision} + \text{recall}}.$$

The β is in range $(0, +\infty)$. The greater the β is, the larger the weight the recall has. β can be selected differently based on the situation. If we want to find all TB-positive people and the cost of additional procedures is not that important, then we give the recall the greater weight. If the cost of additional procedures is crucial, then the precision is much more important even at the cost of skipping some TB-positive people.

This article investigated the following beta threshold options: 0.25, 1, 2. The validation set was used for the threshold calculation. The following metrics were investigated: accuracy (A), which is the number of right guesses, as well as precision (P) and recall (R). Since each cardiogram was divided into several images, the following rule was used to determine the label for the entire cardiogram. If at least one cardiogram is labeled as TB-positive, then the entire cardiogram is labeled TB-positive. The metrics that are determined for the entire ECG are called: AC, PC, and RC for accuracy for ECG, precision for ECG, and recall for ECG, respectively. Since one patient can have multiple ECG, the patient score for cardiogram labels was calculated. The patient score (PS) equals to mean accuracy for person ECGs averaged through all people.

Models without augmentation (M1) and with it (M2) were compared using the metrics described above. F_β metric was used to determine which model has a better ratio between recall and precision. F_β, which shows the ratio between ECG precision and ECG recall is shown as FC_β.

Table II, table III, and table IV show the performance of models trained with augmentation and without it. As shown in these tables, the results are quite promising and customizable, depending on the situation. These tables show that augmentation increases the values of F_β and FC_β when β equals to 0.25 and 1, but does not improve the values when β equals to 2. It can be explained by the fact that augmentation parameters were selected to maximize PR AUC value. This value is most correlated with F_β when β equals to 1. As can be seen from the ratio between precision and recall, the main increase in the PR AUC metric was associated with an increase in precision. It explains the increase in performance when β equal to 0.25. Perhaps it is necessary to focus on other metrics during the process of selection of augmentation parameters to increase the value of F_2.

TABLE II. EXPERIMENT RESULTS

β	P		R		F_β	
	M1	*M2*	*M1*	*M2*	*M1*	*M2*
0.25	0.8918	**0.9147**	0.471	**0.55075**	0.617	**0.687**
1	0.756	**0.785**	**0.7911**	0.782	0.773	**0.784**
2	**0.665**	0.6274	0.878	**0.9393**	**0.757**	0.7523

TABLE III. EXPERIMENT RESULTS

β	PC		RC		FC_β	
	M1	*M2*	*M1*	*M2*	*M1*	*M2*
0.25	0.8688	**0.894**	0.5559	**0.6223**	0.678	**0.734**
1	0.726	**0.76**	**0.863**	0.856	0.789	**0.8059**
2	**0.621**	0.5711	0.937	**0.9825**	**0.7475**	0.7223

TABLE IV. EXPERIMENT RESULTS

β	A		A1		PS	
	M1	*M2*	*M1*	*M2*	*M1*	*M2*
0.25	0.843	**0.866**	0.8655	**0.885**	0.8389	**0.8498**
1	0.876	**0.884**	0.8824	**0.895**	0.819	**0.831**
2	**0.849**	0.834	**0.8388**	0.8076	**0.7611**	0.7246

V. INTERPRETATION

Doctors often want to understand why a neural network made a confident decision. Therefore, the interpretation of the neural network is quite useful. Besides, interpretation can be used to identify previously unknown signs of the disease to think through treatment in the future.

Grad-CAM [26] shows quite good results in interpreting neural networks. Grad-CAM uses the gradient information flowing into the last convolutional layer of CNN to understand the impact on each neuron for a decision of interest. Fig. 4 and 5 show the interpretation results on the pretrained model without augmentation and with it. The lighter the points on the image, the more important these points are. Both images represent the same ECG of TB-positive person. As can be seen from the images, the neural network draws attention to the PQST complexes, the distance between the peaks, and also the height of some peaks. Features selected by the network with and without augmentation are almost identical. However, the network with augmentation does not take into consideration some features as one at the end of the 4th row. Perhaps the elimination of some noisy features might help to improve the classification quality.

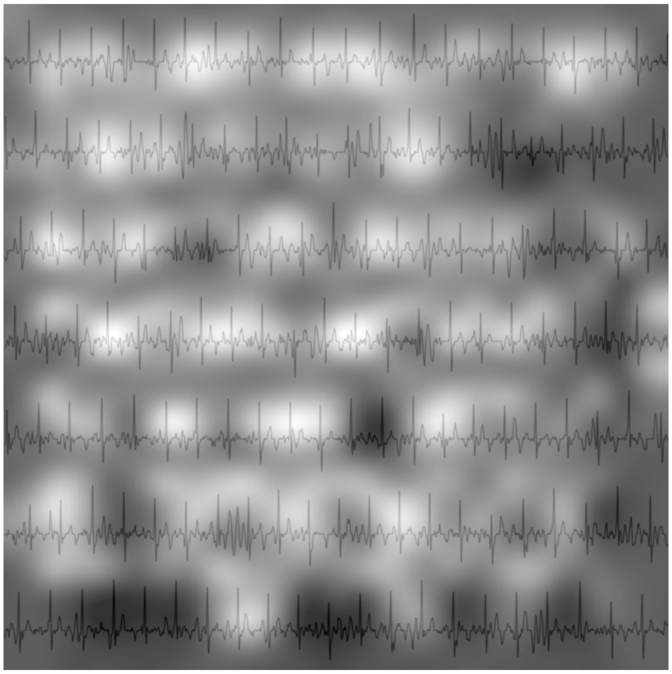

Fig. 4. Interpretation of NN without augmentation on TB-positive person.

VI. Discussion

The reader of this article should know that this article has the following limitations: the sample size of TB-positive people is quite small, so for further investigation of performance, the bigger sample should be used. It should also be noted that the sample contains patients who received treatment for tuberculosis. Due to the existing database structure, it is impossible to say whether the patient took medicine, and if he did, how recent the use is. For more robust experiments, a new database must be compiled. However, even in the current circumstances, this study is useful for the following reasons. Let us imagine an organization that employs a TB-positive person, who has already started taking medicine and hides his condition, thus increasing the chance of infecting people around him. In such a situation, it is still beneficial to screen people with the suggested algorithm. Besides, it is known that people taking medicine may experience an increase in QT-interval [27]. However, in this dataset, the average value of the QT-interval showed a low predictive quality for TB of about 0.51 ROC AUC. Based on this, we can conclude that the neural network finds some other defining signs. Also, with the help of the neural network interpretation in the future, the effect of TB-drugs on the heart can be explored.

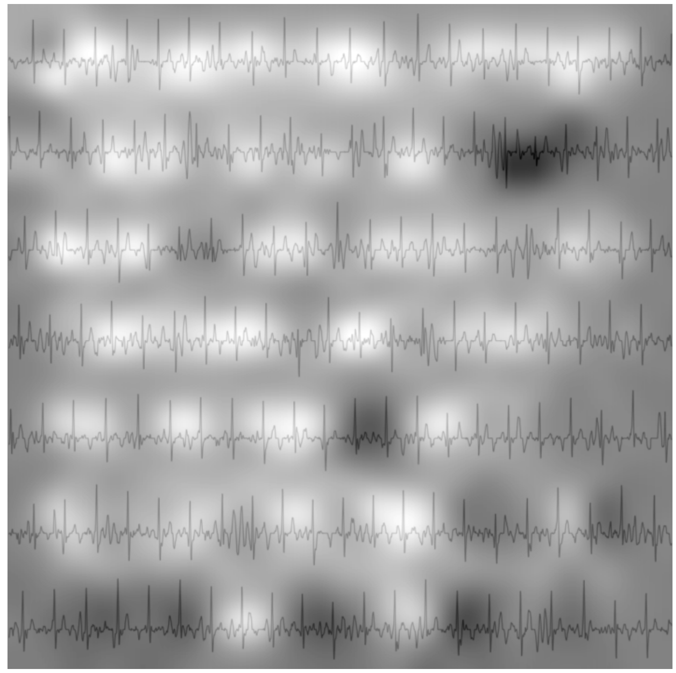

Fig. 5. Interpretation of NN with augmentation on TB-positive person.

VII. Conclusion

The new method for classifying the visualization of ECGs using pre-trained architecture was proposed. This method was used on a single-channel ECG to explore the possibility of determination of tuberculosis from such images. It was determined that the neural network could detect tuberculosis by a cardiogram with moderately high performance. The new method help to improve the quality of the ECG image classification. ROC AUC score was increased from 0.9079 to 0.934, and PR AUC score was increased from 0.7956 to 0.84. The significance of the results shows that this method should be used to classify the ECG for other tasks. For example, to detect coronary heart disease and arrhythmia.

Another paper novelty was the confirmation of the possibility of increasing the classification quality by using online augmentation with such architectures and data. It has been shown that online augmentation can improve the performance of the classification. In this article, the performance was improved from 0.934 ROC AUC to 0.9475 ROC AUC, and PR AUC was increased from 0.84 to 0.875. The increase shows that online augmentation should be used when using the proposed method.

The method for interpretation of such models is described, which can be used further for drug and disease effect exploration purposes.

Various ratios of quality metrics of the obtained solution for tuberculosis detection have been received. They show that it is reasonable to conduct TB pre-screening using mobile ECG in the future to identify potential patients and recommend further specific tests. This study is an important proof of concept.

Some disadvantages of the dataset were indicated in the article. In the future, to build a production solution, it is reasonable to increase and carefully verify the dataset.

Acknowledgment

This article contains the results of a project carried out within the imple-mentation of the Program of the Center for Competence of the National Tech-nology Initiative “Center for Storage and Analysis of Big Data”, supported by the Ministry of Science and Higher Education of the Russian Federation under the Lomonosov Moscow State University Project Support Fund 13/1251/2018 from 11.12.2018.

References

[1] ”Tuberculosis (TB)”, Who.int, 2020 [Online]. URL: https://www.who.int/news-room/fact-sheets/detail/tuberculosis.

[2] ”AliveCor”, Alivecor.com, 2020 [Online]. URL: https://www.alivecor.com.

[3] ”Cardiomonitor (in Russian) CardioQVARK”, Cardioqvark.ru, 2020 [Online]. URL: http://www.cardioqvark.ru.

[4] M. Huaman, D. Henson, E. Ticona, T. Sterling and B. Garvy, ”Tuberculosis and cardiovascular disease: linking the epidemics,” Tropical Diseases, Travel Medicine and Vaccines, vol. 1, no. 1, 2015.

[5] T. Pashina, A. Gaidel, P. Zelter, A. Kapishnikov and A. Nikonorov, ”Automatic highlighting of the region of interest in computed tomography images of the lungs,” Computer Optics, vol. 44, no. 1, pp. 74-81, 2020. DOI: 10.18287/2412-6179-CO-659.

[6] J. Agafonova, A. Gaidel, P. Zelter and A. Kapishnikov, ”Efficiency of machine learning algorithms and convolutional neural network for detection of pathological changes in MR images of the brain,” Computer Optics, vol. 44, no. 2, pp. 266-273, 2020. DOI: 10.18287/2412-6179-CO-671.

[7] S. Kiranyaz, T. Ince and M. Gabbouj, ”Real-Time Patient-Specific ECG Classification by 1-D Convolutional Neural Networks,” IEEE Transactions on Biomedical Engineering, vol. 63, no. 3, pp. 664-675, 2016.

[8] M. Zihlmann, D. Perekrestenko and M. Tschannen, ”Convolutional Recurrent Neural Networks for Electrocardiogram Classification,” Computing in Cardiology Conference (CinC), pp. 1-4, 2017.

[9] Z. Yao, Z. Zhu and Y. Chen, ”Atrial fibrillation detection by multi-scale convolutional neural networks,” 20th International Conference on Information Fusion (Fusion), Xi’an, pp. 1-6, 2017.

[10] R. Banerjee, A. Ghose and S. Khandelwal, ”A Novel Recurrent Neural Network Architecture for Classification of Atrial Fibrillation Using Single-lead ECG,” 27th European Signal Processing Conference (EUSIPCO), A Coruna, Spain, pp. 1-5, 2019.

[11] P. Sean and J. Huizinga, ”Robust and fast heart rate variability analysis of long and noisy electrocardiograms using neural networks and images,” arXiv: 1902.06151, 2019.

[12] M. Kim, H. Ko and S. Pan, ”A study on user recognition using 2D ECG based on ensemble of deep convolutional neural networks,” Journal of Ambient Intelligence and Humanized Computing, vol. 11, no. 5, pp. 1859-1867, 2019.

[13] H. Kaiming, X. Zhang, Sh. Ren and J. Sun, ”Deep residual learning for image recognition,” Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770-778, 2016.

[14] F. Chollet, ”Xception: Deep learning with depthwise convolutions,” Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 1251-1258, 2017.

[15] X. Saining, R. Girshick, P. Dollár, Z. Tu and K. He, ”Aggregated residual transformations for deep neural networks,” Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 1492-1500, 2017.

[16] K. Simonyan, and A. Zisserman, ”Very deep convolutional networks for large-scale image recognition,” arXiv: 1409.1556, 2014.

[17] H. Alaskar, N. Alzhrani, A. Hussain and F. Almarshed, ”The Implementation of Pretrained AlexNet on PCG Classification,” Intelligent Computing Methodologies, pp. 784-794, 2019.

[18] P. Virtanen, "SciPy 1.0: fundamental algorithms for scientific computing in Python," Nature Methods, vol. 17, no. 3, pp. 261-272, 2020.

[19] J. Pan and W. J. Tompkins, "A Real-Time QRS Detection Algorithm," IEEE Transactions on Biomedical Engineering, vol. BME-32, no. 3, pp. 230-236, 1985.

[20] J. Hunter, "Matplotlib: A 2D Graphics Environment", Computing in Science & Engineering, vol. 9, no. 3, pp. 90-95, 2007.

[21] D. Mahajan, R. Girshick, V. Ramanathan, K. He, M. Paluri, Y. Li, A. Bharambe and L. van der Maaten, ”Exploring the limits of weakly supervised pretraining,” Proceedings of the European Conference on Computer Vision (ECCV), pp. 181-196, 2018.

[22] A. Mikołajczyk and M. Grochowski, ”Data augmentation for improving deep learning in image classification problem,” International Interdisciplinary PhD Workshop (IIPhDW), Swinoujście, pp. 117-122, 2018.

[23] V. Guryanova, ”Online Augmentation for Quality Improvement of Neural Networks for Classification of Single-Channel Electrocardiograms,” Communications in Computer and Information Science, pp. 37-49, 2020.

[24] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan and T. Killeen, ”PyTorch: An imperative style, high-performance deep learning library,” Advances in Neural Information Processing Systems, pp. 8024-8035, 2019.

[25] S. van der Walt, S. Colbert and G. Varoquaux, ”The NumPy Array: A Structure for Efficient Numerical Computation,” Computing in Science & Engineering, vol. 13, no. 2, pp. 22-30, 2011.

[26] R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh and D. Batra, ”Grad-cam: Visual explanations from deep networks via gradient-based localization,” Proceedings of the IEEE international conference on computer vision, pp. 618-626, 2017.

[27] A. A. Bykova, M. I. Chashkina, M. V. Serova, A. V. Simonov and F. Yu. Kopylov, "The prevalence of prolongation of the QT interval in patients receiving antituberculous chemotherapy," Cardiology and Cardiovascular Surgery, vol. 12, no. 2, pp. 146-151, 2019.