Pathologies Prediction on Short ECG Signals with Focus on Feature Extraction Based on Beat Morphology and Image Deformation

Jeffrey van Prehn[1], Svetoslav Ivanov[1], Georgi Nalbantov[1]

[1]Data Science Consulting Ltd, Pernik, Bulgaria

Abstract

Introduction. Automated detection of several key cardiac pathologies in reduced-lead ECGs is an enabling factor in applying ECG analysis on a larger scale. The PhysioNet/Computing in Cardiology Challenge 2021 (hereafter referred to as The Challenge) identifies a set of key cardiac pathologies and challenges us with the task to automatically detect them. Critical to this task is the extraction of features from these ECGs which, combined, mark the presence of one or more of these key cardiac pathologies.

Methodology. Algorithms were devised to automatically extract features based on the definitions as used in medical practice, beat morphology and image deformation. A binary classifier for each key cardiac pathology was trained using these features, extracted from the labeled ECGs from The Challenge. The binary classifiers were combined into a multi-label classifier by learning thresholds on the scores of the binary classifiers using Bayesian optimization in a cross-validation setting.

Results. Our contribution submitted for validation achieved a challenge metric score of 0.38, 0.38, 0.38, 0.38 and 0.37 placing us (team DSC) 44, 41, 44, 44 and 44 out of 58 teams on 12-lead, 6-lead, 4-lead, 3-lead and 2-lead validation datasets respectively. Cross validation on the publicly available training data provided by The Challenge resulted in a challenge metric score of 0.32, 0.31, 0.31, 0.32 and 0.32. Our best model had a cross validation challenge metric score of 0.79 on the training set.

1. Introduction

In The PhysioNet/Computing in Cardiology Challenge 2020 [1, 2], participants were challenged to develop open-source algorithms to automatically identify cardiac abnormalities in 12-lead ECG recordings.

Increasing popularity of wearable health monitoring technology paves the way to the application of ECG analysis on a larger scale [3].

However, continuous monitoring of the 12-lead ECG is impractical and unattractive for a wearable system, due to the obtrusiveness and discomfort that the placement and connection of 10 electrodes would cause [4].

The PhysioNet/Computing in Cardiology Challenge 2021 [1,5] focused on automated, open-source approaches for classifying cardiac abnormalities from reduced-lead ECGs.

To this end, several databases are provided to the challenge participants: the CPSC2018 database [6], the CPSC2018-Extra database [2, 5], the INCART database [7], the PTB database [8], the PTB-XL database [9], the Chapman-Shaoxing Database [10], the Ningbo Database [11] and the Georgia database [2, 5]. For a description of these databases as wel as for a description of the databases that were used for validation and testing, and therefore not disclosed, see [2, 5].

Our best entry in The Challenge uses features based on beat morphology and image deformation to train a binary one-versus-rest classifier for each cardiac abnormality.

2. Methods

For The 2021 Challenge we made a contribution based on our contribution to The 2020 Challenge as described in [12]. Furthermore, we developed a model based on image deformation.

The Chapman-Shaoxing and the Ningbo Database were excluded for training in order to reduce training time and because we lacked time to look at them in detail.

2020 contribution. In summary, our contribution to the 2020 Challenge uses features available online, supplemented by features developed by us, which are specific to the pathologies provided in The 2020 Challenge. Some features are based on morphology, others are based on significant ECG points. Using these features, binary classifiers were trained for the scored pathologies and for frequently occurring combinations of pathologies. These classifiers were then combined in a hierarchical and parallel way. Thresholds on the scores output by the classifiers were chosen such that the distribution of predicted labels in a test set was the same as the observed one in the training set.

Participation in the 2020 challenge led to the observa-

Feature / Class	AF	AFL	BBB	Brady	LBBB	RBBB	IRBBB	LAD	LAnFB	LPR	LQRSV	LQT	NSIVCB	NSR	PAC	PR	PVC	RAD	Stach	Tab	Tinv	Leads
Mean RR-interval length	x	x		x										x		x			x			I
Standard deviation of RR-interval length	x	x												x		x						I
Mean length of P waves														x	x		x					I
Distance between Q and S			x		x	x	x						x	x								I
Distance between start of P and Q	x													x								I
Morphology of the signal														x	x		x					A
Distance between start and peak of P and R	x									x				x								I, II
Distance between start of P and Q														x								I, II
Distance between end of P and R										x				x								I, II
Distance between end of T and beginning of next P														x								I, II
Distance between Q and S			x		x	x	x						x	x								I, II
Heart rate				x										x					x			I, II
Heart rate standard deviation	x	x		x										x		x			x			I, II
Difference between amplitude at Q and S								x	x					x				x				I
Distance between P and Q														x	x		x					I, II
Distance between Q and end of T												x		x								I, II
Inverted T (y/n)			x		x	x	x						x	x							x	A (I, II)
RMSE of a linear fit of the T-wave														x						x		A (I, II)
Difference between amplitude at Q and R								x	x					x				x				I-III, aVF
Amplitude of R compared to Q and S											x			x								aV*, V* (I, II)

Table 1: For each class, the features used to identify it and the leads from which the features are determined. 'A' indicates that all available leads are used. For the features related to the T-wave, leads I and II are used for the detection of the T-wave location. All leads are used to detect if the T-wave is inverted. This case is indicated in the table with (I, II). All features are represented by a single number, except the morphology of the signal, which is represented by 20 numbers.

tion that there is room for improvement in the robustness of the model, because there is a large variation in challenge metric score over different combinations of training- and test-sets. The challenge metric score was 0.616 on the validation set and 0.194 on the test set.

The 2020 contribution uses 24 general features and an additional 217 features per used lead.

Minimal model. We improve the robustness of the model by reducing the number of features and by reducing the number of classifiers.

We reduce the number of features by using, for a particular binary classifier, only those features which we deem as relevant for that particular binary classification problem (see Table 1).

The minimal model uses 2 general features. The number of features for the minimal model differs per binary classifier as shown in the table. The total number of features generated in the minimal model is 33, 36, 40, 49 and 73 for 2, 3, 4, 6 and 12-lead ECG's respectively.

Feature selection is done manually by including those features which are used in medical practice.

We reduce the number of classifiers by only training one-versus-rest classifiers for the scored pathologies and not for frequently occurring combinations of pathologies.

Minimal, bayesopt. Like in [13] we now take into account the challenge metric by optimizing the thresholds using the MATLAB® [14] implementation of Bayesian Optimization [15].

Minimal, DS-specific. In [16] it was noted that great variability (in score) between datasets is observed. Therefore, we trained dataset-specific classifiers for the datasets we know have instances in both the training data and the validation data. This results in three classifiers: one on all datasets combined, one on CPSC2018 and CPSC2018-Extra combined and one on the Georgia dataset.

Image deformation. This approach is based on scoring each beat (the ECG signal defined by 3 consecutive R peaks) within the patient's ECG strip. First the signal is rescaled to 250 Hz and possible trend in the signal baseline is removed using a median filter. Then for each available lead the R peaks are estimated using the 'gqrs' [17] program. The R peak location estimations are combined in a universal strip prediction. Each beat is extracted and scaled to 60 beats per minute (500 points). A simple estimator of beat morphology is implemented to find the possible location of the QRS interval, the T-wave and the P-wave. A database

of validated 'obvious' beats for several pathologies was extracted after manual review using a proprietary tool for ECG signal labeling (written in MATLAB® [14]). We call these validated beats atlases. We assume that each atlas is a truthful representation of a labeled pathology. The image deformation distance to a random sample of 10 atlases per pathology is calculated for each new beat. To capture variation, the distances (Euclidean distance, cosine distance, Pearson's correlation distance) of each beat to the mean beat of the strip is computed. Distances are only computed for lead I, in order not to exceed the restrictions on running time. To include information regarding the signal itself we compute the discrete one dimensional wavelet transform of the data and collect the approximation coefficients. We use the Daubechies 1 wavelet [18]. The final training data set is comprised of 291 features. 4 based on Beats per minute, 6 based on distance of a beat to the mean beat of the strip, 180 based in image deformation, 37 based on beat morphology, and 64 based on the discrete wavelet transform. To account for multi-label data, we split each strip with more than one label in two separate instances/strips with individual labels. We favored XGBoost [19] as a machine learning tool because its implementation allowed for easy supply of weights on each observation. We weigh each strip by the reciprocal of the amount of pathologies in this strip to preserve the prior distributions. In order to deal with Unscored pathologies we combine them into a single pathology. Once the XGBoost CV procedure is finished for each pathology we take the mean predicted score for each beat across the strip. This produces our score on strip level. For each pathology we automatically select the score threshold based on the minimum of a simple cost function - 5*FN+1*FP. Each pathology is left with its own score threshold. The cross validation AUC for the chosen thresholds for each pathology is shown in Table 4. No grid search for hyper parameters is run due to computational, human and time constraints. The key XGBoost hyper parameters used were: objective = multi:softprob, max depth = 8, estimators = 200.

3. Results

Table 2 shows the validation scores for our submissions during the official phase of The Challenge.

Table 3 shows the cross-validation scores on the training data for our submissions during the official phase of The Challenge and for the models which we did not manage to submit on time.

Table 4 shows the challenge scores when we leave one dataset out for training and use that dataset for testing. As noted before, the Ningbo and Shaoxing databases are only used for testing.

Table 5 shows the area under the ROC Curve with cross-validation on the training set.

The image deformation model achieved a score of 0.77 with cross-validation on the training set. Only two models were build one for 12 lead and one for 2 lead strips. The number of leads affects only the calibrated R peaks estimation and the beat extraction. The features for each beat are derived only on one lead (lead I).

Model / # leads	12	6	4	3	2
2020 contribution	0.551	0.529	0.538	0.535	0.524
Minimal model	0.385	0.375	0.378	0.383	0.365

Table 2: Challenge metric scores on validation set.

Model / # leads	12	6	4	3	2
2020 contribution	0.532	0.501	0.506	0.502	0.496
2020, bayesopt	0.523	0.495	0.493	0.490	0.492
Minimal model	0.315	0.309	0.313	0.315	0.315
Minimal, bayesopt	0.375	0.366	0.364	0.368	0.359
Minimal, DS-specific	0.238	0.235	0.241	0.241	0.243
Image deformation*	0.772				0.758

Table 3: Challenge metric scores with cross-validation on training set

*The CV dataset for Image deformation consists of 3500 strips selected at random from the dataset with a constraint of minimum 80 examples of a pathology

Dataset / Model	2020 contribution	Minimal, bayesopt
CPSC	0.289	0.299
PTB	-0.356	-0.169
Georgia	0.246	0.241
Ningbo	0.408	0.335
Shaoxing	0.542	0.428

Table 4: Challenge metric scores with cross-validation with leave-one-dataset-out on 12-lead ECG's.

4. Discussion and Conclusions

We implemented a model which uses a very small beat morphology feature set. It achieves a better score on some datasets and a worse score on other. Using additional features based on image deformation greatly improves the score.

References

[1] Goldberger AL, Amaral LA, Glass L, Hausdorff JM, Ivanov PC, Mark RG, Mietus JE, Moody GB, Peng CK, Stanley HE. PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals. Circulation 2000;101(23):e215–e220.

[2] Perez Alday EA, Gu A, Shah A, Robichaux C, Wong AKI, Liu C, Liu F, Rad BA, Elola A, Seyedi S, Li Q, Sharma A, Clifford GD, Reyna MA. Classification of 12-lead ECGs:

Class	2020 contribution	Minimal, bayesopt	Difference	Image deformation*
AF	0,98	0,97	0,01	0,99
AFL	0,95	0,82	0,13	0,99
BBB	0,50	0,50	0,00	0,99
Brady	0,97	0,93	0,05	0,78
LBBB	0,98	0,98	0,00	1,00
RBBB	0,98	0,89	0,09	0,95
IAVB	0,94	0,93	0,01	0,98
IRBBB	0,92	0,76	0,16	0,95
LAD	0,94	0,87	0,07	0,94
LAnFB	0,97	0,91	0,06	0,97
LPR	0,95	0,88	0,07	0,98
LQRSV	0,92	0,83	0,10	0,98
LQT	0,94	0,84	0,10	0,98
NSIVCB	0,79	0,75	0,04	0,96
NSR	0,94	0,90	0,05	0,91
PAC	0,90	0,77	0,13	0,96
PR	1,00	0,80	0,19	0,98
PVC	0,84	0,66	0,18	0,97
Qab	0,78	0,61	0,17	0,96
RAD	0,97	0,89	0,08	0,99
SA	0,96	0,93	0,03	0,98
SB	0,99	0,96	0,03	0,99
Stach	0,99	0,99	0,01	1,00
Tab	0,87	0,73	0,13	0,96
Tinv	0,83	0,65	0,18	0,97

Table 5: Area under the ROC Curve with cross-validation on the 12-lead ECG's in the training set.
*The CV dataset for Image deformation consists of 3500 strips selected at random from the dataset with a constraint of minimum 80 examples of a pathology

the PhysioNet/Computing in Cardiology Challenge 2020. Physiological Measurement 2020;41.

[3] Duncker D, Ding WY, Etheridge S, Noseworthy PA, Veltmann C, Yao X, Bunch TJ, Gupta D. Smart Wearables for Cardiac Monitoring—Real-World Use beyond Atrial Fibrillation. Sensors 2021;21(7).

[4] Figueiredo C, Mendes P. Towards Wearable and Continuous 12-Lead Electrocardiogram Monitoring: Synthesis of the 12-Lead Electrocardiogram Using 3 Wireless Single-Lead Sensors. BIODEVICES 2012 Proceedings of the International Conference on Biomedical Electronics and Devices 01 2012;329–332.

[5] Reyna MA, Sadr N, Perez Alday EA, Gu A, Shah A, Robichaux C, Rad BA, Elola A, Seyedi S, Ansari S, Li Q, Sharma A, Clifford GD. Will Two Do? Varying Dimensions in Electrocardiography: the PhysioNet/Computing in Cardiology Challenge 2021. Computing in Cardiology 2021;48:1–4.

[6] Liu F, Liu C, Zhao L, Zhang X, Wu X, Xu X, Liu Y, Ma C, Wei S, He Z, Li J, Kwee ENY. An Open Access Database for Evaluating the Algorithms of Electrocardiogram Rhythm and Morphology Abnormality Detection. Journal of Medical Imaging and Health Informatics 2018; 8(7):1368—1373.

[7] Tihonenko V, Khaustov A, Ivanov S, Rivin A, Yakushenko E. St Petersburg INCART 12-lead Arrhythmia Database. PhysioBank PhysioToolkit and PhysioNet 2008;.

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

[9] Wagner P, Strodthoff N, Bousseljot RD, Kreiseler D, Lunze FI, Samek W, Schaeffter T. PTB-XL, a Large Publicly Available Electrocardiography Dataset. Scientific Data 2020;7(1):1–15.

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

[11] Zheng J, Cui H, Struppa D, Zhang J, Yacoub SM, El-Askary H, Chang A, Ehwerhemuepha L, Abudayyeh I, Barrett A, Fu G, Yao H, Li D, Guo H, Rakovski C. Optimal Multi-Stage Arrhythmia Classification Approach. Scientific Data 2020;10(2898):1–17.

[12] Nalbantov G, Ivanov S, Van Prehn J. Multi-Class Classification of Pathologies Found on Short ECG Signals. In 2020 Computing in Cardiology. 2020; 1–4.

[13] Vicar T, Hejc J, Novotna P, Ronzhina M, Janousek O. ECG Abnormalities Recognition Using Convolutional Network with Global Skip Connections and Custom Loss Function. In 2020 Computing in Cardiology. 2020; 1–4.

[14] MATLAB. version 9.10 (R2021a). Natick, Massachusetts: The MathWorks Inc., 2021.

[15] Snoek J, Larochelle H, Adams RP. Practical Bayesian Optimization of Machine Learning Algorithms. In Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 2, NIPS'12. Red Hook, NY, USA: Curran Associates Inc., 2012; 2951–2959.

[16] Uguz DU, Berief F, Leonhardt S, Antink CH. Classification of 12-lead ECGs Using Gradient Boosting on Features Acquired with Domain-Specific and Domain-Agnostic Methods. In 2020 Computing in Cardiology. 2020; 1–4.

[17] Silva I, Moody G. An Open-source Toolbox for Analysing and Processing PhysioNet Databases in MATLAB and Octave. Journal of Open Research Software 09 2014;2:e27.

[18] Daubechies I. Ten Lectures on Wavelets. SIAM, 1992.

[19] Chen T, Guestrin C. XGBoost: A Scalable Tree Boosting System. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16. New York, NY, USA: Association for Computing Machinery, 2016; 785–794.

Address for correspondence:

Jeffrey van Prehn
Lippestraat 24, Voerendaal, The Netherlands
jvanprehn@gmail.com