Improving ECG-Based Biometric Identification Using End-to-End Convolutional Networks

João Ribeiro Pinto[1]
jtpinto@fe.up.pt
Jaime S. Cardoso[1]
jaime.cardoso@fe.up.pt
André Lourenço[2]
arl@cardio-id.com

[1] INESC TEC & Faculdade de Engenharia
Universidade do Porto
Porto, Portugal
[2] CardioID Technologies, Instituto Superior de Engenharia de Lisboa, & Instituto de Telecomunicações
Lisboa, Portugal

Abstract

Using Convolutional Neural Networks (CNNs), this work studies how the integration of all processes needed for biometric recognition on a single model improves ECG-based subject identification. An end-to-end unidimensional CNN is proposed, which receives raw blindly-segmented ECG signals and outputs an identification, after being optimised, as a whole, during training. The proposed method was evaluated on the UofTDB collection, offering 96.1% identification rate (IDR), and on the PTB database, attaining 98.6% IDR. When compared with implemented state-of-the-art methods and results reported in the literature, the network showed improved performance and enhanced robustness to the increased noise and variability of off-the-person signals, even with larger sets of subjects.

1 Introduction

The electrocardiogram (ECG) is a biosignal that is defying the dominion of face, fingerprint, voice, and iris over research and industry in biometric recognition. The ECG is universal, sufficiently permanent, and easily measurable, with increased comfort if acquired using non-intrusive techniques. Besides this, its hidden nature and inherent liveness information place it as a promising biometric trait.

Research in ECG-based biometrics [6] has been quickly moving from highly clean and controlled medical signals (designated as *on-the-person*) towards more comfortable and realistic settings (*off-the-person*), using fewer electrodes on the subjects' fingers or palms. With on-the-person signals, Matta *et al*. [5] used linear discriminant analysis and a nearest neighbour classifier with autocorrelation coefficients from ECG segments, while Brás and Pinho [2] used Kolmogorov-based compression on signals denoised using moving average, notch, and lowpass filters. Eduardo *et al*. [3] trained a deep autoencoder to extract signal representations fed to a nearest neighbour classifier, while Salloum and Kuo [8] proposed a Recurrent Neural Network (RNN) with Long Short-Term Memory and Gated-Recurrent Units.

With off-the-person signals, Lourenço *et al*. [4] used average heartbeats, denoised and normalised in time and amplitude. Pinto *et al*. [7] proposed the extraction of DCT coefficients to be used by a Support Vector Machine classifier. Wieclaw *et al*. [10] fed individual segmented heartbeats to a multilayer perceptron (MLP). Zhang *et al*. [11] extracted bidimensional representations the acquired ECG signals and used a 2D Convolutional Neural Network (CNN) for classification. In general, the inferior results reported with off-the-person signals, relative to those with on-the-person signals, illustrate the effect of increased noise and variability on the identification performance.

We observe that, so far, most electrocardiogram-based methods for biometric recognition are composed of several separate processes, each focused and optimised for a single task. Even the recently proposed deep learning algorithms rely on additional techniques for the denoising and feature extraction alongside convolutional or recurrent networks. However, it is reasonable to assume that the separate optimisation of such processes can limit the achievable performance of the recognition methods. Combining all processes in a single network would enable joint, simultaneous, and coordinated optimisation for better performance and robustness to the enhanced noise and variability of off-the-person signals.

Hence, in this work, we study the use of a convolutional neural network (CNN) for ECG-based biometric identification. The proposed network is end-to-end: it receives raw, blindly-segmented electrocardiogram segments and outputs a predicted identity among all enrolled subjects. Thus, it dismisses separate techniques and assumes control over all processes required for robust recognition, expectedly offering improved identification performance and robustness. Two large databases with on-the-person and off-the-person signals were used for performance evaluation. The identification rate (IDR, or accuracy) results were compared with selected state-of-the-art methods and literature results.

2 Proposed Methodology

As aforementioned, the proposed methodology consists of a unidimensional convolutional neural network that receives a raw ECG segment and outputs a decision on the respective identity. The CNN, as shown in Fig. 1 is composed of four convolutional layers, three max-pooling (MaxPool) layers, and one fully-connected (dense) layer.

The convolutional and pooling layers compose the feature-focused part of the model. The first two convolutional layers have 24 filter banks, while the remaining have 36 filter banks, while all filters' size is 1×5. These enable the network to learn to represent the input signal in the most advantageous way for the identification task at hand. Rectified Linear Unit (ReLU) was used as activation for all convolutional layers. The pooling layers, with size 1×5 and placed between each two consecutive convolutional layers, greatly reduce the dimensionality of the feature maps, reducing processing load during both training and inference, effectively making the model more efficient.

Receiving the flattened feature maps from the last convolutional layer, the fully-connected layer is in charge of classification. This layer is composed of N neurons (where N is the total number of enrolled identities) and will, at each neuron, appropriately weigh each input feature to output expected scores for each identity. Softmax activations are used for a normalised distribution of those scores.

As discussed above, the model is optimised as a whole, for the task at hand, in order to maximise the achievable identification performance. The training of the network is performed using the optimiser Adam, based on sparse categorical cross-entropy loss, with empirically tuned learning rate.

During training, both dropout and data augmentation were used to avoid overfitting. Dropout was used between the last convolutional layer and the fully-connected layer. Unidimensional data augmentation was applied to the training signals by dividing each input segment into five 1 second subsegments and randomly shuffling them. The augmentation was implemented using an online data generator.

3 Evaluation Settings

The proposed methodology was evaluated on the UofTDB [9] and PTB [1] collections. UofTDB is a collection of off-the-person ECG signals acquired at 200 Hz using dry electrodes on the fingertips of 1019 subjects, in up to six sessions (over a period of up to six months) on five different postures. It enables the study of the impact of long-term variability and movement/activity noise on the recognition performance. PTB holds 15-lead ECG signals from 290 subjects at rest, acquired at 1000 Hz in medical settings. For this study, we resampled PTB signals to 200 Hz and, as common in the literature, we used solely Lead I.

The signals were divided into five-second segments (1000 samples) to make up the dataset, which was randomly divided 70% for training and 30% for testing. IDR (identification rate, or classification accuracy) was selected as the performance metric. The state-of-the-art methods of Eduardo *et al*. [3] and Matta *et al*. [5] were implemented and tested in the same conditions. Along with the entire UofTDB dataset of 1019 subjects, two subdatasets with 25 and 100 subjects were used to study performance in smaller populations.

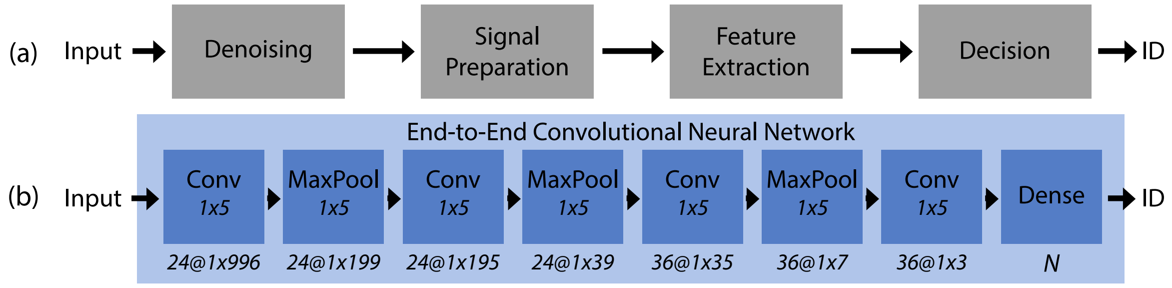

Figure 1: The conventional structure of an ECG-based biometric identification algorithm (a), compared with the proposed end-to-end network (b).

Table 1: Performance evaluation results of the proposed method and the implemented state-of-the-art approaches.

Method	IDR per database and number of subjects			
	UofTDB			PTB
	25	100	1019	290
Proposed Method	99.7%	98.7%	96.1%	98.6%
AC/LDA [5]	96.8%	95.5%	90.0%	98.8%
Autoencoder [3]	97.0%	93.6%	85.0%	99.5%

Table 2: Comparison between the proposed method's performance and the results reported in the literature (N.S. - number of subjects in the database; O.P. - off-the-person).

Method	Database and N.S.	O.P.	IDR
Proposed Method	UofTDB - 1019	Yes	96.1%
Wieclaw *et al.* [10]	Private - 18	Yes	89.0%
Lourenço *et al.* [4]	Private - 16	Yes	94.3%
Zhang *et al.* [11]	Several - 10	Yes	98.4%
Pinto *et al.* [7]	Private - 6	Yes	94.9%
Proposed Method	PTB - 290	No	98.6%
Salloum and Kuo [8]	ECG-ID - 90	No	100%
Brás and Pinho [2]	PTB - 52	No	99.9%

4 Results and Discussion

The performance results for the evaluated methods are presented in Table 1. The proposed network achieved 96.1% when evaluated with the entire UofTDB dataset of 1019 subjects and 98.6% on the PTB dataset. These are, arguably, desirable results for an ECG-based identification system, considering the acquisition settings and the large sets of subjects.

With the on-the-person PTB dataset, the best results were attained by the Autoencoder method proposed by Eduardo *et al.* [3]. However, with the more challenging UofTDB datasets, the proposed method outperformed both state-of-the-art methods in all cases, showing improved robustness to the increased noise and variability carried by off-the-person signals. With these datasets, the IDR difference between the proposed method and the best state-of-the-art method was 2.9% for 25 subjects, but increases to 6.1% for 1019 subjects, showcasing better scalability of the proposed network to larger sets of identities.

When compared with other performance results recently reported in the literature, the proposed method shows promise. Despite the different evaluation settings and much larger number of identities on the evaluation datasets, the proposed method outperformed all literature methods evaluated on off-the-person ECG signals (see Table 2) except the one proposed by Zhang *et al.* [11] and, with on-the-person signals (PTB), it was able to nearly match the almost perfect performance reported on the literature.

5 Conclusion and Future Work

This work focuses on the study of an end-to-end convolutional neural network for biometric identification based on minimally obtrusive, off-the-person electrocardiogram signals. The proposed network was able to adequately integrate all processes needed for recognition in a single model that receives raw signals and outputs predicted identities.

Evaluated with large ECG signal collections, the proposed method was successful in outperforming other state-of-the-art methods with signals acquired in off-the-person settings, and offered competitive performance when compared with several on-the-person results reported in the literature. The network showed improved identification performance, even with larger sets of subjects and considerably increased noise and variability of off-the-person signals, ultimately showing promise for successful implementation in real biometric applications.

Acknowledgements

This work is financed by the ERDF - European Regional Development Fund through the Operational Programme for Competitiveness and Internationalisation - COMPETE 2020 Programme within project "POCI-01-0145-FEDER-006961", and by National Funds through the Portuguese funding agency, FCT - Fundação para a Ciência e a Tecnologia, as part of project "UID/EEA/50014/2013", and within the PhD grant number "SFRH/BD/137720/2018".

The authors wish to express their gratitude to the BioSec.Lab of the University of Toronto for granting access to the UofTDB database, and to acknowledge Physionet and the creators of the PTB database.

References

[1] R. Bousseljot, Kreiseler D., and A. Schnabel. Nutzung der EKG-Signaldatenbank CARDIODAT der PTB über das Internet. *Biomedizinische Technik*, 40(s1):317–318, 1995.

[2] S. Brás and A. J. Pinho. ECG biometric identification: A compression based approach. In *37th EMBC*, pages 5838—-5841, 2015.

[3] A. Eduardo, H. Aidos, and A. Fred. ECG-based Biometrics using a Deep Autoencoder for Feature Learning: An Empirical Study on Transferability. In *ICPRAM 2017*, pages 463–470, 2017.

[4] A. Lourenço, H. Silva, and A Fred. Unveiling the Biometric Potential of Finger-based ECG Signals. *Computational Intelligence and Neuroscience*, 2011:720971, 2011.

[5] R. Matta, J. Lau, F. Agrafioti, and D. Hatzinakos. Real-time continuous identification system using ECG signals. In *24th CCECE*, pages 1313 – 1316, 2011.

[6] J. R. Pinto, J. S. Cardoso, and A. Lourenço. Evolution, Current Challenges, and Future Possibilities in ECG Biometrics. *IEEE Access*, 6:34746–34776, 2018.

[7] J. R. Pinto, J. S. Cardoso, A. Lourenço, and C. Carreiras. Towards a Continuous Biometric System Based on ECG Signals Acquired on the Steering Wheel. *Sensors*, 17(10):2228, 2017.

[8] R. Salloum and C. C. J. Kuo. ECG-based biometrics using recurrent neural networks. In *ICASSP 2017*, pages 2062–2066, 2017. doi: 10.1109/ICASSP.2017.7952519.

[9] S. Wahabi, S. Pouryayevali, S. Hari, and D Hatzinakos. On evaluating ECG biometric systems: session-dependence and body posture. *IEEE TIFS*, 9(11):2002–2013, 2014.

[10] L. Wieclaw, Y. Khoma, P. Falat, D. Sabodashko, and V. Herasymenko. Biometric identification from raw ECG signal using deep learning techniques. In *9th IEEE IDAACS*, pages 129–133, 2017.

[11] Q. Zhang, D. Zhou, and X. Zeng. Pulseprint: Single-arm-ECG biometric human identification using deep learning. In *2017 IEEE UEMCON*, pages 452–456, 2017.