Advancing ECG Biometrics to the Deep Learning Era: Towards a Complete End-to-End Solution

João Ribeiro Pinto

Faculdade de Engenharia da Universidade do Porto
Rua Dr. Roberto Frias, 4200-465 Porto, Portugal
jtpinto@fe.up.pt

Abstract: Like several other pattern recognition fields, electrocardiogram (ECG) biometrics has been quickly moving towards deep learning approaches, in a bid to overcome powerful noise and variability in realistic scenarios. However, traditional processes are still commonly appended to deep learning approaches, possibly limiting achievable performance. This paper presents our pioneering work on truly end-to-end ECG biometrics, addressing the integration of traditional pipeline processes, integrated template security, and interpretability. The results show that the proposed methodology is able to offer improved performance and robustness with challenging realistic data, without overlooking biometric security. As such, it is a complete deep learning solution for ECG biometrics, offering an accurate, robust, and promising foundation for real ECG biometric applications.

1 Introduction

Biometric recognition has been quickly invading every aspect of our lives. And as sensitive applications such as border control and bank accounts grow more reliant on biometrics, the value of successfully spoofing biometric systems becomes more and more enticing. Major biometric characteristics such as face or fingerprints, although offering high accuracies, are continuously endangered by being relatively easy to steal without the victim's awareness. As such, the so-called *medical biometrics*, especially the Electrocardiogram (ECG), have been highly regarded due to their inherent liveness information and hidden nature which creates strong natural obstacles against spoofing [EPR+17, PCL18].

Two decades ago, researchers presented the earliest proofs-of-concept for biometric algorithms based on ECG signals [BPPW01]. These first approaches relied on fiducial features: time, amplitude, and slope measurements between specific ECG heartbeat landmarks. These features were easy to compute for medical signals, which are cleaner and more complete but require subjects to be at rest with several electrodes attached to their bodies (including the chest). However, these acquisition configurations are obviously incompatible with realistic biometric applications.

As such, ECG biometrics has been steadily evolving towards off-the-person settings: simpler acquisition settings, with fewer electrodes with minimally intrusive placements, allowing for free movements, and overall focused on higher user comfort and acceptability [PCL18]. This has resulted in enhanced noise and variability which ultimately turned fiducial approaches unreliable. Holistic approaches [ABH12] and, more recently, deep learning models [EAF17] have enabled higher performances but are still far behind those reported for face or fingerprints. However, the field had yet to delve into truly end-to-end deep learning models as well as other relevant topics on the path to real applicability, such as template security or model interpretability.

This paper overviews our research effort to advance ECG biometrics towards more accurate, complete, and robust real applications through deep learning. Our contributions focus on three main aspects: (a) studying the integration of the typical ECG biometrics pipeline into a simple but accurate end-to-end CNN architecture for improved performance; (b) proposing the Secure Triplet Loss to promote cancelability and unlinkability on end-to-end biometric models, without needing separate protection processes; and (c) studying the behaviour of our ECG biometric approach using explainability tools, to analyse its decisions and understand the relative importance of each signal region.

The two latter topics were both explored using the end-to-end model from the first topic as a framework. This results in the overall contribution of a cohesive and comprehensive study on deep-learning-powered ECG biometrics, from performance and optimisation to template security and explainability.

2 Methodology

End-to-end deep learning Our work advancing the topic of ECG biometrics towards the deep learning era began in [PCL19] with the proposal of an end-to-end model for biometric identification, later adapted for identity verification in [PC19]. Our convolutional neural network (CNN) model (see Fig. 1) has one first part for signal processing and feature extraction (composed of four convolutional layers), and one last part for decision/classification (composed of fully connected layers). It receives blindly-segmented five-second lead I ECG signals and outputs either identity scores (for identification) or a template (for identity verification).

We studied the progressive integration of typical pipeline processes (signal denoising, preparation, feature extraction, and decision) into the deep learning model. This aimed to assess the benefits of a holistically-optimised end-to-end model, trying to obtain higher robustness to the extra noise and variability of off-the-person data. Additionally, we com-

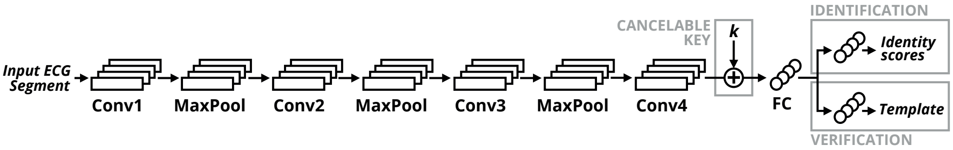

Figure 1: The architecture of the proposed end-to-end solution for ECG biometrics.

pared the use of cross-entropy identification learning *vs.* triplet loss learning.

Ensuring template security To truly achieve robust commercial ECG biometric applications, biometric security is essential. However, one should not have to let go of the performance advantages brought by truly end-to-end deep learning approaches. As such, we developed the Secure Triplet Loss (STL) [PCC21], a simple strategy that allows deep biometric models to simultaneously learn both biometric recognition and security. Unlike other template protection methods [GBRG$^+$16, DBR$^+$19], it does not regard security as an afterthought but promotes it from the start within the biometric model itself.

Inspired by the triplet loss, the STL (1) trains a model to transform biometric samples x into templates y whose distance d is smaller if they share the same identity. The model also receives a generated security key k, concatenated with the latent features before the first fully-connected layer, that is bound with x (see Fig. 1). The STL leads the model to learn cancelability by also promoting larger d if the respective k's do not match. Unlinkability is learned by promoting a low Kullback-Leibler divergence (D_{KL}) between the different-key distance distributions when identities match ($P_{d_{DP}}$) or not ($P_{d_{DN}}$). Thus, the distance between y's can be used to accept/reject an identity claim, and changing k quickly invalidates compromised templates, ensuring biometric security.

$$l = \gamma \max\left[0, \alpha + d_{SP} - \min(\{d_{SN}, d_{DP}, d_{DN}\})\right] + (1-\gamma) D_{KL}(P_{d_{DP}} || P_{d_{DN}}) \quad (1)$$

Explaining decisions Decisions made by deep learning models are notoriously difficult to explain, as these models are typically highly sophisticated unlike simpler and more explainable models such as decision trees. As ECG biometrics evolves toward the deep learning era, the transparency of identity decisions becomes more and more important. In this work [PC20], we study the identification decisions of our end-to-end methodology using explainability tools: Occlusion [ZF14], Saliency [SVZ14], Gradient SHAP [LL17], and DeepLIFT [SGK17]. We explore cleaner *vs.* noisier data and growing sets of identities to evaluate the model's behaviour in diverse scenarios. Beyond applying interpretability to ECG biometrics, we aimed to study the relative importance of heartbeat waveforms and the literature claim that the QRS complex alone is enough to recognise identity.

3 Results

End-to-end deep learning Our model was evaluated on the UofTDB [WPHH14], a challenging off-the-person database with data from 1019 identities. As visible in Table 1, using the end-to-end model brings considerable performance improvements when compared to a traditional ECG biometrics pipeline. The flexibility of the model allows it to take full advantage of the information of raw signals and learn to overlook the noise and variability present in off-the-person ECG data.

This is also visible in the identity verification results in Table 2. The proposed end-to-end model offers the best results, considerably better than the state-of-the-art alternatives.

Table 1: Evolution of identification accuracy results on the UofTDB off-the-person database when integrating further traditional pipeline stages into a deep learning model [PCL19].

Configuration	**Accuracy**
Signal Processing + Segmentation + Feature Extraction + CNN	90.6%
Signal Processing + Segmentation + CNN	93.1%
Signal Processing + CNN	94.2%
End-to-end CNN	96.1%

Identification training attains better performance, but the distance to triplet loss training is reduced when in more challenging scenarios with fewer enrollment data.

Table 2: Identity verification equal error rate (EER) results with triplet loss (TL) *vs*. identification learning (IT) with growing enrollment data duration, compared with the state-of-the-art [PC19].

	Enrollment Duration			
Method	**5s**	**10s**	**15s**	**30s**
IT	13.70%	10.92%	9.52%	7.86%
TL	13.93%	11.89%	10.90%	9.94%
AC/LDA [ABH12]	30.27%	17.90%	16.55%	15.82%
Autoencoder [EAF17]	21.82%	19.68%	18.84%	17.09%
DCT [PCL19]	23.05%	20.41%	18.55%	17.38%

Ensuring template security Security was evaluated on two parameters: cancelability and linkability. Cancelability is measured by the rate of false matches in instances where the keys don't match at the EER point ($FMR_C@EER$) and linkability is measured using the $D_{\leftrightarrow}^{sys}$ proposed by [GBRG$^+$16]. As presented in Table 3, the proposed approach offers perfect cancelability and very good linkability results at the expense of slightly degraded performance *vs*. the original triplet loss. It also outperforms the state-of-the-art approach Bloom Filters (BF) [GBRG$^+$16] in all considered measures and Homomorphic Encryption (HE) [DBR$^+$19] in cancelability. It was also significantly less complex and, thus, faster than HE in our experiments, achieving the goal of template security within the end-to-end model.

Table 3: Summary of the results on identity verification, cancelability, and linkability with the proposed Secure Triplet Loss, the original triplet loss, and the state-of-the-art [PCC21].

	Performance	**Cancelability**	**Linkability**
Method	$EER\,(\%)$	$FMR_C@EER$	$D_{\leftrightarrow}^{sys}$
Secure Triplet Loss	13.58	0.0	0.005
Triplet Loss	12.56	-	-
BF [GBRG$^+$16]	15.76	0.0075	0.234
HE [DBR$^+$19]	12.49	0.0806	0.002

Explaining decisions The identification model was also evaluated using explainability tools to understand its behaviour regarding the relative importance of the different parts of the ECG signal. Here, the PTB database [BKS95], with on-the-person data from 290 identities, was used alongside UofTDB. Results on either database and with varying sets of subjects are presented in Fig. 2. One can notice that, for less challenging conditions, especially with on-the-person data, the QRS complex is the focus of almost all attention. However, as we move towards larger identity sets and off-the-person data, the model feels the need to use information from all parts of the ECG heartbeats. This indicates that, for truly accurate and robust approaches in challenging realistic scenarios, the practice of using just the QRS complex should be discarded.

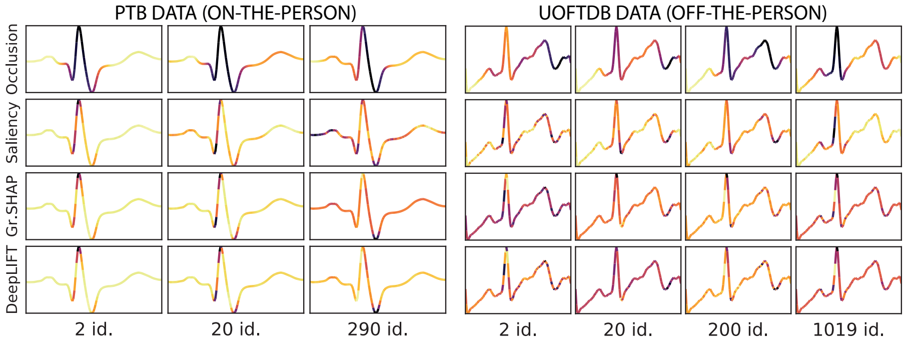

Figure 2: Average heartbeat explanations for PTB and UofTDB data with growing number of identities (darker regions denote higher relevance for the decision as measured by each method).

4 Conclusion

This work focused on advancing the topic of ECG biometrics by proposing a complete end-to-end methodology with integrated template security and explainability studies. Our goal was to take full advantage of the capabilities of deep learning and build an integrated solution for real ECG biometric applications.

Performance-wise, this work shows the benefits of having an end-to-end approach versus a traditional pipeline of separate processes in challenging off-the-person scenarios. Security-wise, we show that it is possible to learn biometric security in a simple and integrated way within end-to-end models. Finally, experiments on interpretability offer deeper insights into the relative importance of ECG waveforms for identification. Although further efforts should be devoted to this topic, especially on the creation of more complete and realistic public databases, the solution proposed here offers a solid and promising foundation for real ECG biometric applications.

References

[ABH12] F. Agrafioti, F. M. Bui, and D. Hatzinakos. Secure Telemedicine: Biometrics for Remote and Continuous Patient Verification. *Journal of Computer Networks and Communications*, page 11, 2012.

[BKS95] R. Bousseljot, D. Kreiseler, and A. Schnabel. Nutzung der EKG-Signaldatenbank CARDIODAT der PTB über das Internet. *Biomedizinische Technik*, 40(s1), 1995.

[BPPW01] L. Biel, O. Pettersson, L. Philipson, and P. Wide. ECG analysis: a new approach in human identification. *IEEE Transactions on Instrumentation and Measurement*, 50(3):808–812, Jun 2001.

[DBR^{+}19] P. Drozdowski, N. Buchmann, C. Rathgeb, M. Margraf, and C. Busch. On the Application of Homomorphic Encryption to Face Identification. In *2019 International Conference of the Biometrics Special Interest Group (BIOSIG)*, 2019.

[EAF17] A. Eduardo, H. Aidos, and A. Fred. ECG-based Biometrics using a Deep Autoencoder for Feature Learning: An Empirical Study on Transferability. In *International Conference on Pattern Recognition Applications and Methods (ICPRAM)*, pages 463–470, Porto, Portugal, Feb. 2017.

[EPR^{+}17] S. Eberz, N. Paoletti, M. Roeschlin, M. Kwiatkowska, and I. Martinovic. Broken Hearted: How To Attack ECG Biometrics. In *Network and Distributed System Security Symposium*, 2017.

[GBRG^{+}16] M. Gomez-Barrero, C. Rathgeb, J. Galbally, C. Busch, and J. Fierrez. Unlinkable and irreversible biometric template protection based on bloom filters. *Information Sciences*, 370-371:18–32, 2016.

[LL17] S. M. Lundberg and S.-I. Lee. A Unified Approach to Interpreting Model Predictions. In *Advances in Neural Information Processing Systems*, pages 4765–4774, 2017.

[PC19] J. R. Pinto and J. S. Cardoso. An End-to-End Convolutional Neural Network for ECG-Based Biometric Authentication. In *10th International Conference on Biometrics Theory, Applications and Systems (BTAS)*, Tampa, FL, United States, Sep. 2019.

[PC20] J. R. Pinto and J. S. Cardoso. Explaining ECG Biometrics: Is It All In The QRS? In *2020 International Conference of the Biometrics Special Interest Group (BIOSIG)*, Darmstadt, Germany, Sep. 2020.

[PCC21] J. R. Pinto, M. V. Correia, and J. S. Cardoso. Secure Triplet Loss: Achieving Cancelability and Non-Linkability in End-to-End Deep Biometrics. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 3(2):180–189, 2021.

[PCL18] J. R. Pinto, J. S. Cardoso, and A. Lourenço. Evolution, Current Challenges, and Future Possibilities in ECG Biometrics. *IEEE Access*, 6:34746–34776, 2018.

[PCL19] J. R. Pinto, J. S. Cardoso, and A. Lourenço. Deep Neural Networks For Biometric Identification Based On Non-Intrusive ECG Acquisitions. In *The Biometric Computing: Recognition and Registration*, chapter 11, pages 217–234. 2019.

[SGK17] A. Shrikumar, P. Greenside, and A. Kundaje. Learning Important Features through Propagating Activation Differences. In *34th International Conference on Machine Learning*, pages 3145–3153, 2017.

[SVZ14] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps. In *International Conference on Learning Representations Workshops (ICLR)*, 2014.

[WPHH14] S. Wahabi, S. Pouryayevali, S. Hari, and D. Hatzinakos. On Evaluating ECG Biometric Systems: Session-Dependence and Body Posture. *IEEE Transactions on Information Forensics and Security*, 9(11):2002–2013, Nov 2014.

[ZF14] M. D. Zeiler and R. Fergus. Visualizing and Understanding Convolutional Networks. In *Computer Vision – ECCV 2014*, pages 818–833, 2014.