Secure Triplet Loss for End-to-End Deep Biometrics

João Ribeiro Pinto, Jaime S. Cardoso, Miguel V. Correia

INESC TEC & FEUP (joao.t.pinto@inesctec.pt)

IWBF 2020 - April 30th, 2020

Context & Motivation

Context & Motivation

The Performance-Security Duality

Biometric recognition is everywhere.
And users are demanding...

... very high performance:

- Accuracy;
- Robustness;
- Speed;
- Lightweight.

... and data protection:

- Irreversibility;
- Cancelability;
- Non-linkability;
- Robustness to attacks.

...but current literature works focus mostly on one side.

Context & Motivation

The Performance-Security Duality

So, methods are either...

... too focused on performance...

- sophisticated deep learning methods;
- very high accuracy and robustness;
- poor template security;
- and/or wide performance gap.

... or too focused on security.

- predesigned feature extraction methods;
- cancelability based on biohashing or encryption methods;
- subpar performance.

Why don't we have both?

Context & Motivation

Goals & Contributions

Deep learning has been able to learn so many difficult things.
Why not template security?

Goal:
Take advantage of the properties of end-to-end deep learning to **achieve secure templates *and* improved performance.**

Contributions:
- A novel triplet loss formulation for secure biometric templates;
- A strategy for the inclusion of cancelability keys on end-to-end models;
- First secure end-to-end deep method for ECG biometrics;
- A thorough evaluation of performance and security.

Secure Triplet Loss

Triplet Loss

Original Formulation

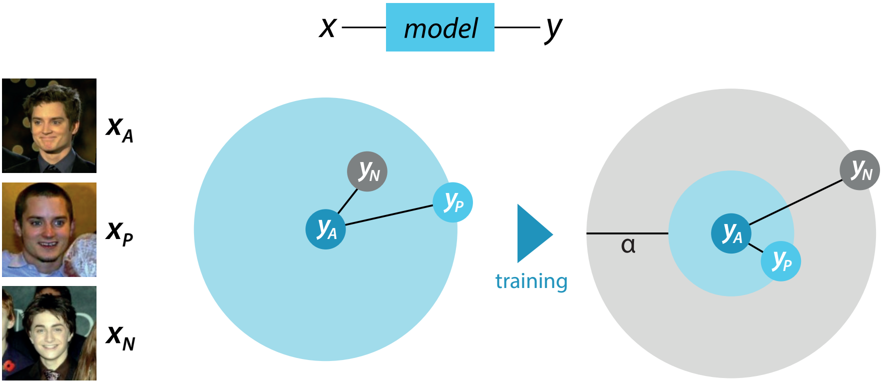

Quite good for performance. But not for cancelability...

Pictures of Elijah Wood and Daniel Radcliffe from LFW Face Database: `http://vis-www.cs.umass.edu/lfw/`.

Secure Triplet Loss

Learning Cancelability

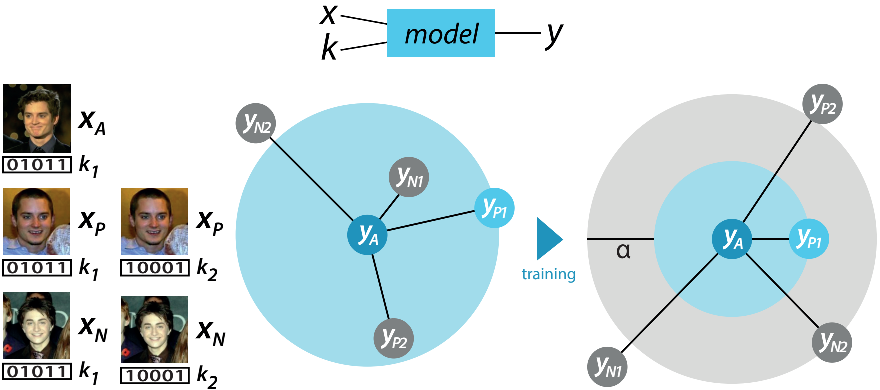

Pictures of Elijah Wood and Daniel Radcliffe from LFW Face Database: `http://vis-www.cs.umass.edu/lfw/`.

Secure Triplet Loss

Triplet Loss:

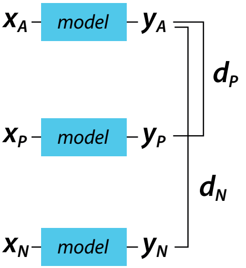

$$l = \max\left[0, \alpha + d_P - d_N\right]$$

Secure Triplet Loss:

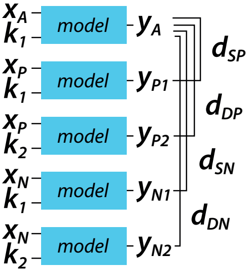

$$l = \max\left[0, \alpha + d_{SP} - \min(\{d_{SN}, d_{DP}, d_{DN}\})\right]$$

Experiments and Results

Experiments and Results

Model and Training

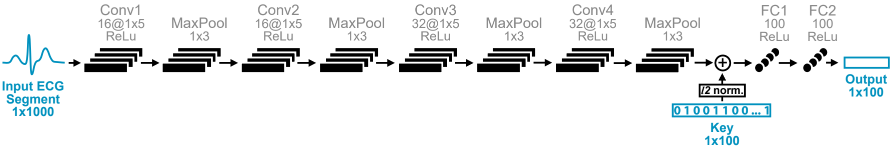

Data from the **UofTDB**[1] off-the-person ECG database:

- Each sample is a blind 5s recording segment (at F_s=200Hz);
- Data from 100 subjects for training:
  90 000 triplets generated for training, 10 000 for validation;
- Data from 918 subjects for testing:
  10 000 triplets generated.

[1] Wahabi *et al.*, "On Evaluating ECG Biometric Systems: Session-Dependence and Body Posture", *IEEE TIFS*, 2014.

Experiments and Results

Performance

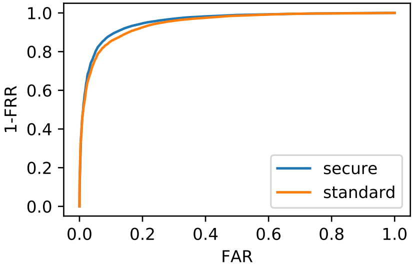

10.63% EER *vs.* 12.55% with the original loss

Better than the state-of-the-art in off-the-person ECG biometrics[1,2].

[1] Pinto *et al.*, "An End-to-End Convolutional Neural Network for ECG-Based Biometric Authentication", *BTAS*, 2019.

[2] Pinto *et al.*, "Evolution, Current Challenges, and Future Possibilities in ECG Biometrics", *IEEE Access*, 2018.

Experiments and Results
Cancelability

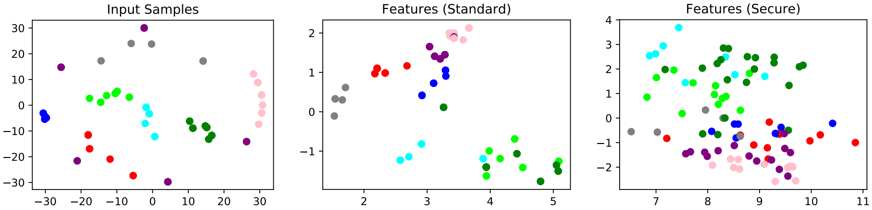

With the Secure Triplet Loss, the model does not cluster samples by identity when keys don't match.

Experiments and Results

Cancelability

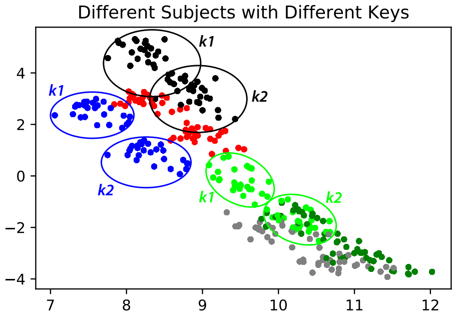

But when keys do match, samples are neatly clustered by identity.
These figures also show the behaviour observed when changing keys.

Experiments and Results

Non-linkability

Evaluation based on the $D_{\leftrightarrow}(s)$ and $D_{\leftrightarrow}^{sys}$ linkability measures[1].

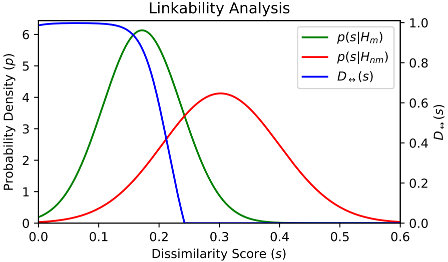

$D_{\leftrightarrow}^{sys}$=0.67 (between semi- and fully-linkable)

[1] Gomez-Barrero *et al.*, "Unlinkable and irreversible biometric template protection based on bloom filters", *Information Sciences*, 2016.

Experiments and Results

Other Security Measures

	Original	*Secure*
Privacy Leakage Rate[1]	0	0
Secrecy Leakage[1]	-	0
Secret Key Rate[1]	14.20 bits	103.73 bits

The perfect *PLR* and *SL* scores probably result from the properties of end-to-end neural networks[2], which are highly beneficial for secure biometrics.

[1] Using Paul Brodersen's Entropy Estimator for Python: `https://github.com/paulbrodersen/entropyestimators`.

[2] Tishby and Zaslavsky, "Deep learning and the information bottleneck principle", *IEEE ITW*, 2015.

Conclusion

Conclusion

- We can indeed have high performance and template security;
- The Secure Triplet Loss achieves that in a simple way;
- Biometric performance gap is closed;
- Cancelability and irreversibility are ensured;
- Only drawback is high linkability.

Future work:

- Adapt the loss to enforce non-linkability;
- Explore for other biometric traits;
- Explore for different key binding strategies;
- Devise a new triplet mining technique.

Thank you!

Questions? Contact me at *joao.t.pinto@inesctec.pt*.

Acknowledgements:
This work was financed by the ERDF – European Regional Development Fund through the Operational Programme for Competitiveness and Internationalization - COMPETE 2020 Programme and by National Funds through the Portuguese funding agency, FCT – Fundação para a Ciência e a Tecnologia within project "POCI-01-0145-FEDER-030707", and within the PhD grant "SFRH/BD/137720/2018". The authors wish to acknowledge the creators and administrators of the UofTDB database (University of Toronto, Canada), which was essential for this work.

INESCTEC U.PORTO FEUP FACULDADE DE ENGENHARIA UNIVERSIDADE DO PORTO COMPETE 2020 PORTUGAL 2020 UNIÃO EUROPEIA Fundo Europeu de Desenvolvimento Regional FCT Fundação para a Ciência e a Tecnologia