A. Brömme, C. Busch, A. Dantcheva, C. Rathgeb and A. Uhl (Eds.): BIOSIG 2018,
Lecture Notes in Informatics (LNI), Gesellschaft für Informatik, Bonn 2018 1

Towards Fingerprint Presentation Attack Detection Based on Convolutional Neural Networks and Short Wave Infrared Imaging

Ruben Tolosana[1], Marta Gomez-Barrero[2], Jascha Kolberg[2], Aythami Morales[1], Christoph Busch[2], Javier Ortega-Garcia[1]

Abstract: Biometric recognition offers many advantages over traditional authentication methods, but they are also vulnerable to, for instance, presentation attacks. These refer to the presentation of artifacts, such as facial pictures or gummy fingers, to the biometric capture device, with the aim of impersonating another person or to avoid being recognised. As such, they challenge the security of biometric systems and must be prevented. In this paper, we present a new fingerprint presentation attack detection method based on convolutional neural networks and multi-spectral images extracted from the finger in the short wave infrared spectrum. The experimental evaluation, carried out on an initial small database but comprising different materials for the fabrication of the artifacts and including unknown attacks for testing, shows promising results: all samples were correctly classified.

Keywords: presentation attack detection, biometrics, fingerprint, SWIR, CNN

1 Introduction

Deep Learning (DL) has become a thriving topic in the last years [GBC16], allowing computers to learn from experience and understand the world in terms of a hierarchy of simpler units. This way, DL has enabled significant advances in complex domains such as natural language processing [SVL14] and computer vision [Zh16], among many others. The main reasons to understand its high deployment lie on the increasing amount of available data, which thereby allows the succesfull training of deep architectures. These can in turn outperform other traditional machine learning techniques. However, the belief that DL architectures can be only used for those tasks with massive amounts of available data is changing thanks to the development of, for instance, pre-trained models. This concept refers to network models that are trained for a given task with large available databases, and then are retrained (a.k.a. fine-tuned, adapted) for a different task for which data are usually scarce. All these advances have allowed the deployment of DL architectures in many different fields, such as biometric recognition [RD17, To18].

Biometrics refers to automated recognition of individuals based on their biological (e.g., iris or fingerprint) or behavioural (e.g., signature or voice) characteristics. Even if biometric recognition systems offer numerous advantages over traditional authentication methods

[1] BiDA Lab - Biometrics and Data Pattern Analytics, Universidad Autonoma de Madrid, Spain, {ruben.tolosana, aythami.morales, javier.ortega}@uam.es

[2] da/sec - Biometrics and Internet Security Research Group, Hochschule Darmstadt, Germany, {marta.gomez-barrero, jascha.kolberg, christoph.biusch}@h-da.de

(e.g., they provide a stronger link between subject and identity, and they cannot be lost or forgotten), they are also vulnerable to external attacks. Among all possible attack points [IS16], the biometric capture device is probably the most exposed one: no further knowledge about the inner functioning of the system is required to launch an attack. Such attacks are known in the ISO/IEC IS 30107 [IS16] as *presentation attacks* (PA), and refer to the presentation to the capture device of a *presentation attack instrument* (PAI), such as a fingerprint overlay, in order to interfere with its intended behaviour.

To be able to prevent PAs, techniques able to automatically distinguish between bona fide (i.e., real or live) presentations and access attempts carried out by means of PAIs must be developed [MNL14]. They are known as *presentation attack detection* (PAD) methods. A considerable attention has been directed to the development of efficient PAD approaches within the last decade for several biometric characteristics, including iris [GGB17], fingerprint [SB14], or face [GMF14]. In particular, within the DL community, Convolutional Neural Networks (CNNs) have been used for fingerprint PAD purposes, based either on the complete fingerprint samples [NdALM16, Ja17] or on a patch-wise manner [TCB17, CCJ18]. In addition, a Deep Belief Network (DBN) system with multiple layers of Restricted Boltzmann Machines (RBMs) was used in [Ki16] also for fingerprint PAD. A more general approach was tested on face, iris, and fingerprint data in [Me15].

All the aforementioned DL PAD approaches achieve accuracy detection rates over 90% on the freely available LivDet [Li17] and ATVS-FFp databases [Ga11]. Such high accuracy rates indicate not only the valuable contributions of the authors but also that other databases, comprising a larger number of materials for the fabrication of the PAIs, should be explored. However, one question remains unanswered: will unknown attacks also be detected? From the aforementioned works, only Chugh *et al.* considered a wider database comprising over twelve different PAI fabrication materials in [CCJ18]. Part of the materials were used as unknown attacks, showing that the error rates were multiplied up to six times with respect to the evaluation carried out on known attacks. Therefore, some more research is needed in this area.

To further tackle these issues with unknown attacks, some researchers have started considering other sources of information different of the traditional fingerprint capture devices [GGB17, SB14]. More specifically, the use of multi-spectral infrared technologies has been studied for face [St16] and fingerprint [RNB08, Ch11]. More recently, the characteristic remission properties of the human skin for multi-spectral Short Wave Infrared (SWIR) wavelengths was exploited in [St16] for facial PAD, achieving a 99% detection accuracy.

In this context, we propose a fingerprint PAD method based on CNNs and multi-spectral SWIR finger samples captured in the range 1200 nm – 1550 nm. To the best of our knowledge, this is the first study exploring the potential of SWIR imaging and CNNs for fingerprint PAD, on a small database in terms of samples but considering a wide variety of PAI species (i.e., both complete thick gummy fingers and more challenging overlays). We successfully detected all of them. It should also be highlighted that only six PAIs were used for training, thus being able to test the remaining six PAIs as *unknown attacks* (i.e., attacks not seen previously by the classifier, thereby representing a bigger challenge and a better representation of a real-world scenario).

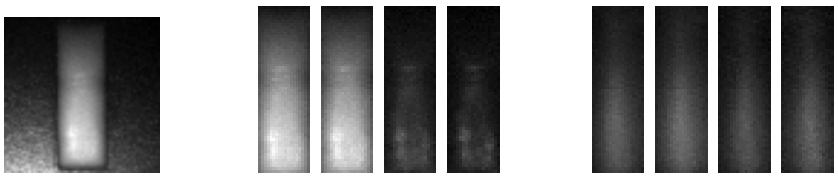

(a) (b) (c) (d) (e) (f) (g) (h) (i)

Figure 1: (a) Bona fide sample captured at 1200 nm, and cropped ROIs corresponding to: (b) to (e) a bona fide sample, and (f) to (i) a silicone PAI, for the selected wavelengths.

The rest of the article is organised as follows. The SWIR sensor and fingerprint PAD method proposed are described in Sect. 2. Sect. 3 presents the experimental protocol and the results obtained in this work. Final conclusions are drawn in Sect. 4.

2 Proposed Presentation Attack Detection Method

2.1 Short Wave Infrared (SWIR) Imaging

The capture device comprises two sensors for SWIR and visible spectrum (VIS) wavelenghts, which are placed next to each other inside a closed box. Next to them, the LEDs for the corresponding wavelengths illuminate the finger. The box includes an open slot on the top where the user stands the finger during the acquisition. When the finger is placed there, all ambient light is blocked and thus only the desired wavelengths are used for the acquisition. In particular, we have used a Hamamatsu InGaAs SWIR sensor array, which captures images of 64×64 pixels, with a 25 mm fixed focal length lens optimised for wavelengths within 900 – 1700 nm. We have considered the following SWIR wavelengths: 1200 nm, 1300 nm, 1450 nm, and 1550 nm, similar to the ones considered in [St16] for the skin vs. non-skin facial classification. Fig. 1a shows the acquired image of a bona fide sample for the 1200 nm wavelength. In addition, and although is out of the scope of this work, fingerprint verification can be carried out with contactless finger photos acquired in the visible spectrum with a 1.3 MP camera and a 35 mm VIS-NIR lens.

In order to utilise only foreground finger information, a preprocessing stage is first applied to the original image (Fig. 1a) so as to select the region of interest (ROI), corresponding to the open slot where the finger is placed. The ROIs of the bona fide sample for all SWIR wavelengths, with a size of 18×58 px, are depicted from Fig. 1b to 1e.

Finally, Fig. 1f to 1i also shows a silicone PAI. Some differences may be observed if we compare the images to those captured from a bona fide presentation: whereas for the bona fide, the images show a decrease in the intensity value for bigger wavelengths, this is not the case for the PAI. Such trend will be hence exploited by the PAD method.

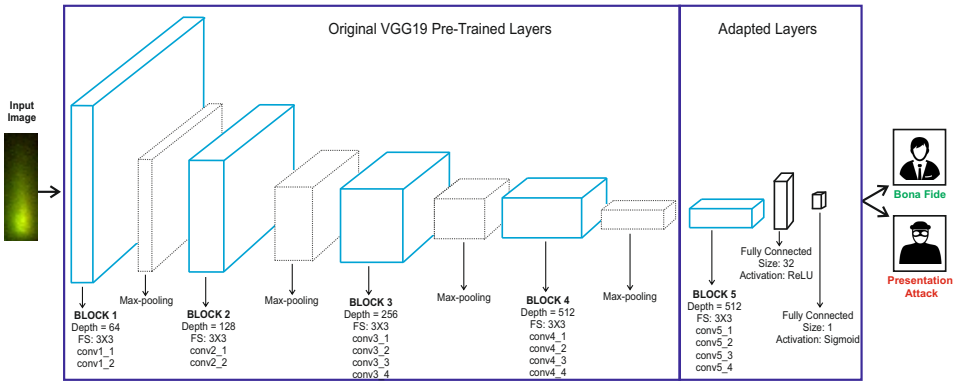

Figure 2: Architecture of our Proposed CNN-based system for fingerprint PAD. FS denotes the filter size of the kernels.

2.2 Convolutional Neural Networks (CNNs)

CNNs have been one of the most successful network architectures in the last years. Some of their key design principles were drawn from the findings of the Neurophysiologists Nobel Prizes David Hubel and Torsten Wiesel in the field of human vision [GBC16]. CNN-based systems are mainly composed of convolutional and pooling layers. The former extracts patterns from the images through the application of several convolutions in parallel to local regions of the images. These convolutional operations are carried out by means of different kernels (adapted by the learning algorithm) that assign a weight to each pixel of the local region of the image depending on the type of patterns to be extracted. Therefore, each kernel of one convolutional layer is focused on extracting different patterns such as horizontal or vertical edges. The output of these operations produces a set of linear activations (a.k.a. feature map) that serve as input to nonlinear activations such as the rectified linear activation function (ReLU). Finally, it is common to use pooling layers to make the representation invariant to small translations of the input. The pooling function replaces the output of the net at a certain region with a statistical summary of the nearby outputs. For instance, the max-pooling function selects the maximum value of the region.

Since, to the best of our knowledge, there are no public databases of SWIR finger images, the available data is not enough to train the entire CNN from scratch. Therefore, we propose a combination of CNN pre-trained models and fine-tuning. Fine-tuning techniques have a two-fold objective, namely: i) replace and retrain the classifier (i.e., the fully-connected layers) of the pre-trained model to our specific task, and ii) adapt the weights of all or some of the convolutional layers. In particular, we have used the VGG19 pre-trained model [SZ15], which achieved the second place in the classification task of the ImageNet 2014 challenge with a total of 1,000 classes such as animals, vehicles, etc. This model comprises a total of 16 convolutional layers and 3 fully-connected layers, and has been modified for the specific task of fingerprint PAD.

Fig. 2 shows the final architecture of the proposed system. The input of the network is an RGB image where each dimension consists of SWIR images captured at different wave-

lengths or combinations of them. In order to optimise the input, an exhaustive analysis was carried out using a development dataset to minimise the intra-class variability of the bona fide class, and at the same time maximise the inter class variability between bona fide and PA samples. The best combination found was: *i*) 1550 nm, *ii*) 1450 nm, and *iii*) a combination of both wavelengths.

Then, given the data scarcity and the small input size, we have reduced the complexity of the original VGG19 pre-trained model by eliminating one of the 3 fully-connected layers. In addition, the number of neurons of the first fully-connected layer is reduced to 32 instead of 512. Regarding the retraining of the VGG19 model, as depicted in Fig. 2, the first 4 convolutional blocks of the CNN network are frozen, whereas the weights of the last convolutional block and fully-connected layers are adapted to the fingerprint PAD task (see vertical line separating the two groups of layers). The reason behind this fine-tuning technique lies on the fact that the first layers of the CNN extract more general features related to directional edges and colours, whereas last layers of the network are in charge of extracting more abstract features related to the specific task.

Finally, the softmax classification layer of the original VGG19 pre-trained model is replaced for a sigmoid activation layer in order to provide output scores between 0 (bona fides) and 100 (PAs), as required in the ISO/IEC 30107-3 on PAD evaluation and reporting [IS17].

3 Experimental Evaluation

The proposed CNN-based system is implemented under the Keras framework using Tensorflow as back-end, with a NVIDIA GeForce GTX 1080 GPU. For the fine-tuning of the layers we consider Adam optimizer with a learning rate value of 0.001 and a loss function based on binary cross-entropy. In the next sections, we describe the experimental protocol followed and discuss the results obtained.

3.1 Database and Experimental Protocol

The selection of the PAI fabrication materials is based on the requirements of IARPA ODIN program evaluation, covering the most challenging PAIs. In particular, the following twelve different PAIs are considered in the experiments: 3D printed fingerprint and 3D printed fingerprint coated with silver paint to mimic the conductive properties of the skin; fingers fabricated with blue and green wax, gelatine, playdoh, silly putty, and silicone; overlays fabricated with dragon skin and urethane; and fingerprints printed on regular matte paper and on transparency paper. The bona fide samples have been captured from seven out of ten fingers in order to increase the variability. For each bona fide and PAI, between one and four samples have been acquired. This database was captured by our BATL project partners at the University of South California.

In order to perform a clear analysis of our proposed approach, the database has been divided into development and test datasets. Moreover, the development dataset is divided

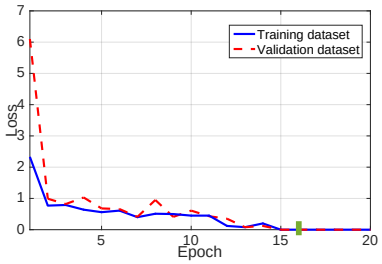

Figure 3: Evolution of the loss function with the number of epochs for the development datasets. Our selected CNN model is indicated using a small green vertical line.

into training and validation datasets that are used for selecting the best configuration of our proposed CNN-based system and adjusting the weights of the final layers of the network. The training dataset comprises 6 bona fides and 6 PAI samples. For the PAIs, we have chosen one sample of dragon skin, blue wax, gelatine, playdoh, silicon, and printed fingerprint. Regarding the validation dataset, a total of 3 bona fides and 3 PAI samples are considered. In this case, only blue wax, playdoh and printed fingerprint PAIs are selected. Finally, the test dataset considered for the final evaluation of the system includes the samples not seen for the network during the development stage (4 bona fides and 38 PAs). It is important to remark that only six out of twelve available PAIs are used for the development of the proposed method in order to evaluate the robustness of our system to new types of PAIs that can arise (i.e., unknown attacks).

Finally, in compliance with the ISO/IEC IS 30107-3 on Biometric presentation attack detection - Part 3: Testing and Reporting [IS17], the following metrics are used to evaluate the performance of the PAD method: *i*) Attack Presentation Classification Error Rate (APCER), or percentage of attack presentations wrongly classified as bona fide presentations; and *ii*) Bona Fide Presentation Classification Error Rate (BPCER), or percentage of bona fide presentations wrongly classified as presentation attacks.

3.2 Results

Fig. 3 shows the evolution of the loss function with the number of epochs for the development datasets. It is important to remark that very similar loss values are obtained for both training and validation datasets along all the epochs, thus showing the robustness of the features extracted by the CNN. The proposed CNN model is selected after 16 epochs, providing a final loss value of 0 for both training and validation datasets.

Then, the fingerprint PAD method is evaluated using new samples from the test dataset. It is worth noting that these samples have not been used during the development of the system, thus yielding unbiased results. The proposed approach achieves final values of 0% APCER and BPCER, proving the success of considering fine-tuning techniques over pre-trained model even with small amounts of data.

Finally, it should be highlighted that all new types of PAIs (e.g. green wax, urethane, or silly putty), which were not considered during the development of our proposed CNN-based system, are correctly detected as PAs. This means that the classifier is robust to all previously unseen attacks considered, thereby proving the soundness of the approach. We may thus conclude that the proposed SWIR sensor and fingerprint PAD method are able to detect unknown attacks.

4 Conclusions

We have presented a novel fingerprint presentation attack detection approach based on CNNs and SWIR multi-spectral images. Based on an exhaustive analysis of the intra- and inter-class variability, two SWIR wavelengths and their combination were selected as input for the network.

The experimental evaluation yields a BPCER = 0% (i.e., highly convenient system) and at the same time APCER = 0% (i.e, highly secure). In fact, even unknown attacks are correctly detected, thereby showing the promising performance of the proposed method, in spite of the small training set (six bona fides and six PAIs). This is partly due to the use of pre-trained CNN models.

As future work lines, we will acquire a bigger database, comprising more PAIs and more bona fide samples, in order to further test the performance of the algorithm for both known and unknown attacks.

Acknowledgements

This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA) under contract number 2017-17020200005. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein. This work was carried out during an internship of R. Tolosana at da/sec. R. Tolosana is supported by a FPU Fellowship from Spanish MECD.

References

[CCJ18] Chugh, T.; Cao, K.; Jain, A.-K.: Fingerprint Spoof Buster: Use of Minutiae-Centered Patches. IEEE TIFS, 13(9):2190–2202, 2018.

[Ch11] Chang, S; Larin, K. et al.: Fingerprint spoof detection by NIR optical analysis. In: State of the Art in Biometrics, pp. 57–84. InTech, 2011.

[Ga11] Galbally, J.; Fierrez, J. et al.: Evaluation of direct attacks to fingerprint verification systems. Telecommunication Systems, 47(3-4):243–254, 2011.

[GBC16] Goodfellow, I.; Bengio, Y.; Courville, A.: Deep Learning. MIT Press, 2016.

[GGB17] Galbally, J.; Gomez-Barrero, M.: Presentation Attack Detection in Iris Recognition. In (Busch, C.; Rathgeb, C., eds): Iris and Periocular Biometrics. IET, August 2017.

[GMF14] Galbally, J.; Marcel, S.; Fierrez, J.: Biometric antispoofing methods: A survey in face recognition. IEEE Access, 2:1530–1552, 2014.

[IS16] ISO/IEC JTC1 SC37 Biometrics: . ISO/IEC 30107-1. Information Technology - Biometric presentation attack detection - Part 1: Framework. ISO, 2016.

[IS17] ISO/IEC JTC1 SC37 Biometrics: . ISO/IEC IS 30107-3. Information Technology - Biometric presentation attack detection - Part 3: Testing and Reporting. ISO, 2017.

[Ja17] Jang, H.-U.; Choi, H.-Y. et al.: Fingerprint Spoof Detection Using Contrast Enhancement and Convolutional Neural Networks. In: Proc. ICISA. pp. 331–338, 2017.

[Ki16] Kim, S.; Park, B. et al.: Deep belief network based statistical feature learning for fingerprint liveness detection. Pattern Recognition Letters, 77:58–65, 2016.

[Li17] LivDet - Liveness Detection Competitions, 2009–2017.

[Me15] Menotti, D.; Chiachia, G. et al.: Deep representations for iris, face, and fingerprint spoofing detection. IEEE TIFS, 10(4):864–879, 2015.

[MNL14] Marcel, S.; Nixon, M.; Li, S.-Z., eds. Handbook of Biometric Anti-Spoofing. Springer, 2014.

[NdALM16] Nogueira, R.-F.; de Alencar Lotufo, R.; Machado, R. C.: Fingerprint Liveness Detection Using Convolutional Neural Networks. IEEE TIFS, 11(6):1206–1213, 2016.

[RD17] Rattani, A.; Derakhshani, R.: On Fine-Tuning Convolutional Neural Networks for Smartphone Based Ocular Recognition. In: Proc. IJCB. 2017.

[RNB08] Rowe, R.-K.; Nixon, K.-A.; Butler, P.-W.: Multispectral Fingerprint Image Acquisition. In (Ratha, N.-K.; Govindaraju, V, eds): Advances in Biometrics: Sensors, Algorithms and Systems. Springer London, pp. 3–23, 2008.

[SB14] Sousedik, C.; Busch, C.: Presentation attack detection methods for fingerprint recognition systems: a survey. IET Biometrics, 3(4):219–233, 2014.

[St16] Steiner, H.; Sporrer, S et al.: Design of an active multispectral SWIR camera system for skin detection and face verification. Journal of Sensors, 2016, 2016.

[SVL14] Sutskever, I.; Vinyals, O.; Le, Q.-V.: Sequence to Sequence Learning with Neural Networks. In: Proc. NIPS. 2014.

[SZ15] Simonyan, K.; Zisserman, A.: Very Deep Convolutional Networks for Large-Scale Image Recognition. In: Proc. ICLR. 2015.

[TCB17] Toosi, A.; Cumani, S.; Bottino, A.: CNN Patch-Based Voting for Fingerprint Liveness Detection. In: Proc. IJCCI. 2017.

[To18] Tolosana, R.; Vera-Rodriguez, R. et al.: Exploring Recurrent Neural Networks for On-Line Handwritten Signature Biometrics. IEEE Access, pp. 1 – 11, 2018.

[Zh16] Zhou, B.; Khosla, A. et al.: Learning Deep Features for Discriminative Localization. In: Proc. CVPR. 2016.