Towards Fingerprint Presentation Attack Detection Based on Short Wave Infrared Imaging and Spectral Signatures

Marta Gomez-Barrero*, Jascha Kolberg, and Christoph Busch

da/sec - Biometrics and Internet Security Research Group,
Hochschule Darmstadt, Germany
email: {marta.gomez-barrero, jascha.kolberg, christoph.busch}@h-da.de

Abstract

Biometric verification systems are currently being deployed in numerous large-scale and everyday applications. Among other vulnerabilities, presentation attacks directed to the sensor (e.g., using a face mask or a gummy finger) pose a severe security threat. To prevent such attacks, presentation attack detection (PAD) techniques have been proposed in the last decade. For the particular case of fingerprint recognition, most approaches are based on conventional optical or capacitive sensors, acquiring a single image, and which can thus detect only a limited number of materials used to fool the sensor.

In this paper we propose a PAD algorithm based on normalised spectral signatures extracted from Short Wave InfraRed (SWIR) images captured at four different wavelengths. It has been shown that the information contained in the selected SWIR wavelengths can help to discriminate skin from other materials. The extracted features are classified using a Support Vector Machine (SVM), thereby yielding a fast real-time performance which can be implemented on almost any application. The experimental evaluation shows that all but one material considered can be detected, including unknown materials (i.e., not used to train the classifier).

1 Introduction

Biometrics refers to automated recognition of individuals based on their biological (e.g., iris or fingerprint) or behavioural (e.g., signature or voice) characteristics [1]. Certainly, biometric recognition is very attractive and useful for the final user: forget about PINs and passwords, you are your own key [1]. In addition, they provide a stronger link between the subject and the action or event. All these facts have allowed a wide deployment of biometric systems in the last decade for different applications, ranging from border control to smartphone unlocking.

In spite of their numerous advantages, biometric systems are vulnerable to external attacks. Among the different possible points of attack summarised in [2], the biometric capture device is probably the most exposed one: in order to launch an attack on the capture device, no further knowledge about the inner functioning of the system is required. Such attacks are known in the literature as *presentation attacks* and defined within the ISO/IEC 30107 standard on biometric presentation attack detection [3] as the "*presentation to the biometric data capture subsystem with the goal of interfering with the operation of the biometric system*". In other words, an attacker can present the capture device with a *presentation attack instrument* (PAI), such as a gummy finger or a fingerprint overlay, in order to impersonate someone else (i.e., active impostor) or to avoid being recognised due to black-listing (i.e., identity concealer).

In order to prevent such attacks, *presentation attack detection* (PAD) methods have been proposed for several biometric characteristics, including iris [4, 5], fingerprint [6, 7] or face

*The author presented this paper at the NISK 2018 conference.

Proceedings of the 11th Norwegian Information Security Conference, Sep. 19-20, 2018.
NISK Journal, e-ISSN 1894-7735. https://ojs.bibsys.no/index.php/NISK

[8]. This term refers to any technique that is able to automatically distinguish between bona fide (i.e., real or live) presentations and access attempts carried out by means of PAIs [9] and has attracted a considerable attention within the last decade. In fact, several international projects, like the European Tabula Rasa [10] and BEAT [11], or the more recent US Odin research program [12] deal with these security concerns. In addition, the LivDet – liveness detection competition series on iris and fingerprint [13] have been running since 2009.

As any other pattern recognition task, PAD techniques can benefit from data acquired with different sensors, thereby providing complementary sources of information [4, 7]. In particular, the use of multi-spectral technologies have been studied for face [14, 15] and fingerprint [16, 17] in the Near InfraRed (NIR) domain. In addition, Hengfoss *et al.* [18] analysed extensively the multi-spectral signatures of living against the cadaver fingers using spectroscopy techniques for wavelengths between 400 nm and 1630 nm. However, no PAIs were analysed in their work.

In contrast to the NIR wavelengths under 850 nm used in most of the aforementioned articles, it has been shown that human skin shows characteristic remission properties for multi-spectral SWIR wavelengths, which are independent of a person's age, gender or skin type [19]. Based on this principle, Steiner *et al.* described in [15, 20] a new skin detection approach based on the spectral signatures extracted from wavelengths ranging from 935 nm to 1550 nm. The authors trained a system based on a Support Vector Machine (SVM) to discriminate skin vs. non-skin (e.g., hair or make-up) pixels, thereby enabling a fast classification process. In their experiments on facial images partially covered with masks, hair or make-up, they achieved a pixelwise classification accuracy over 99.9%.

Inspired by the skin detection method presented in [15, 20], we propose a PAD method based on the spectral signatures of finger samples captured in the range 1200 nm – 1550 nm. To the best of our knowledge, this is the first fingerprint PAD method exploring SWIR imaging. In particular, pixels are classified as skin or non-skin, and a final decision (i.e., bona fide or presentation attack) for the complete sample is made based on the proportion of non-skin pixels detected. We have tested a wide variety of PAI species including both complete thick gummy fingers and more challenging overlays, fabricated with twelve different materials, and successfully detected all materials but one.

It should be highlighted that only six samples where used for training, thus being able to test the remaining six materials 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). Even if large databases comprising thousands of bona fides are freely available, this is not the case for all PAIs. In particular, the LivDet competitions [13] offer large databases comprising bona fides and PAIs. However, only a few distinct PAIs are considered. Therefore, being able to train the system with a small number of samples, stemming from a reduced number of PAIs, is crucial to be able to detect unknown attacks. The pixel-wise classification of the proposed method allows the use of such a small training set.

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

2 Proposed Presentation Attack Detection Method

We describe in this section both the capture device and the PAD method proposed in detail. In particular, we define the normalised spectral signatures for each pixel and the classification

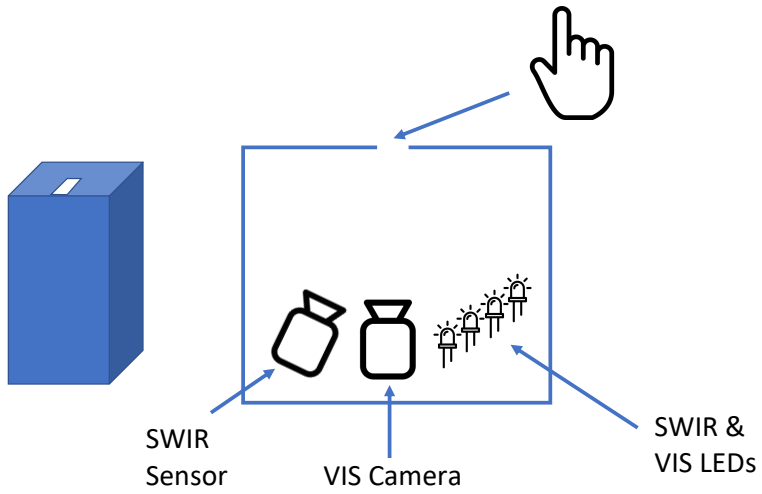

Figure 1: Finger sensor diagram. On the left, the complete box showing the open slot to place the finger for the acquisition. On the right, a diagram of the inner components: two different sensors for the SWIR images and the visible (VIS) light images, together with the corresponding LEDs.

algorithm to reach a final bona fide vs. presentation attack decision for the sample at hand.

2.1 Short Wave InfraRed (SWIR) Imaging

A diagram of the finger SWIR capturing device developed for the present work is included in Fig. 1. As it may be observed, the camera and lens are placed inside a closed box, which includes an open slot on the top. When the finger is placed there, all ambient light is blocked and therefore only the desired wavelengths are used for the acquisition. In particular, we have used a Hamamatsu InGaAs SWIR sensor array, which captures 64×64 px images, with a 25 mm fixed focal length lens optimised for wavelengths within $900 - 1700$ nm. In this article, we have considered the following SWIR wavelengths: 1200 nm, 1300 nm, 1450 nm, and 1550 nm, similar to the wavelengths considered in [15, 20] for the skin vs. non-skin facial classification. An example of the acquired images for a bona fide sample is shown in Fig. 2a for the 1200 nm wavelength. In addition, 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, which are placed next to the SWIR sensor within the closed box (see Fig. 1).

In order to process the images, the central finger-slot region corresponding to the open slot where the finger is placed needs to be extracted from the background. Given that the finger is always placed over the open slot, and the camera does not move, a simple fixed size cropping can be applied. The corresponding bona fide samples for all SWIR wavelengths, with a size of 18×58 px, are depicted in Fig. 2b to 2e.

Finally, samples acquired from three PAIs fabricated with different materials are included in Fig. 3: (a) to (d) a silicone finger, (e) to (h) a dragon skin overlay, and (i) to (l) an orange playdoh finger. Some differences may be observed if we compare the images to those captured from a bona fide presentation in Fig. 2. Whereas for the bona fide, the images show a decrease in the intensity value for bigger wavelengths, this is not the case for the silicone and the dragon skin. Such trend will be hence exploited by the PAD method.

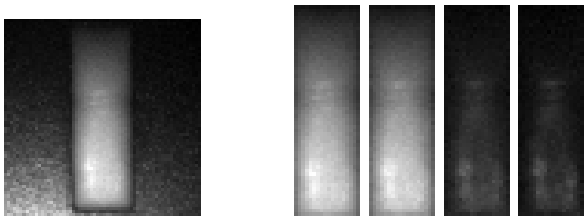

(a) (b) (c) (d) (e)

Figure 2: Bona fide samples: (a) complete image at 1200 nm, and cropped finger-slot regions, where each image was captured at (b) 1200 nm, (c) 1300 nm, (d) 1450 nm, and (e) 1550 nm, respectively.

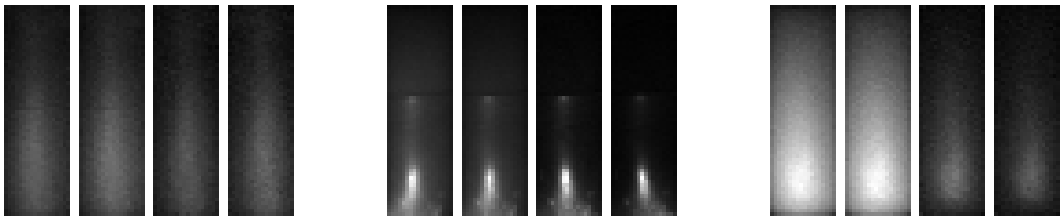

(a) (b) (c) (d) (e) (f) (g) (h) (i) (j) (k) (l)

Figure 3: Cropped finger-slot regions corresponding to several presentation attack instruments: (a) to (d) silicone finger, (e) to (h) dragon skin overlay, and (i) to (l) playdoh finger. In all cases, each image was captured at 1200 nm, 1300 nm, 1450 nm, and 1550 nm, respectively.

We may also see in both Fig. 2 and 3 that the upper and lower ends of the finger present a non-uniform illumination and reflections. For this reason, the PAD algorithm will only analyse the central part of the images, thus yielding a final ROI comprising 18×20 px (see Fig. 4).

2.2 Spectral Signatures

Similar to [15, 20], the spectral signature $\mathbf{ss}$ of a pixel with coordinates (x, y) is defined as:

$$\mathbf{ss}\,(x, y) = (i_1, \ldots, i_N) \tag{1}$$

where i_n represents the intensity value for the nth wavelength. In our particular case study, $N = 4$.

In order to account for illumination changes, and therefore achieve a signature independent of the absolute brightness of the image at hand, a normalised signature is computed. In addition, in order to capture the distinct trends accross different wavelengths shown in Fig. 2 for the bona fides (i.e., skin pixels) and in Fig. 3 for the PAIs (i.e., non-skin pixels), a final normalised difference vector is computed as follows:

$$d\,[i_a, i_b] = \frac{i_a - i_b}{i_a + i_b} \tag{2}$$

with $a, b \leq N$, $a \neq b$ and $-1 \leq d\,[i_a, i_b] \leq 1$. In other words, the normalised differences between all possible wavelength combinations are computed. For our case study with $N = 4$, a total

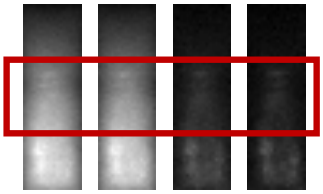

Figure 4: Selected ROI comprising the central 18×20 px from the images acquired at each wavelength.

number of six differences are calculated:

$$\mathbf{d}(x,y) = (d[i_1,i_2], d[i_1,i_3], d[i_1,i_4], d[i_2,i_3], d[i_2,i_4], d[i_3,i_4]) \tag{3}$$

These normalised difference vectors $\mathbf{d}(x,y)$ will be used to classify skin vs. non-skin pixels.

2.3 Final Classification

The final bona fide vs. presentation attack decision for the sample at hand is made in a two step manner:

- For each pixel with coordinates (x,y), the normalised spectral signature $\mathbf{d}(x,y)$ is computed and classified as skin vs. non-skin with a Support Vector Machine (SVM) classifier.
- Given that a single sample comprises a total number of $18 \times 20 = 360$ px per wavelength, the final score s returned by the PAD method will be the proportion of non-skin pixels of the sample ROI in a range of 0 to 100, as per [21].

In other words, regarding the final score $s \in [0, 100]$ generated by the PAD algorithm, low values close to 0 will represent bona fide samples and high values close to 100 will denote presentation attacks.

It should be finally noted that, whereas other high performing state of the art methods need a high number of images for training, up to 1000, the present method, being based on a pixelwise classification, requires a very low number of images for training. In our experiments, only six images per class are utilised.

3 Experimental Evaluation

The PAD algorithm proposed in Sect. 2 is evaluated in this section. First, the database and evaluation metrics are described. We subsequently present and analyse the results obtained.

3.1 Experimental Setup

In the experiments, we have analysed the following twelve different PAIs:

- 3D printed fingerprint and 3D printed fingerprint coated with silver paint (denoted Ag) to mimic the conductive properties of the skin;
- fingers fabricated with blue and green wax, gelatine, playdoh, silly putty and silicone;

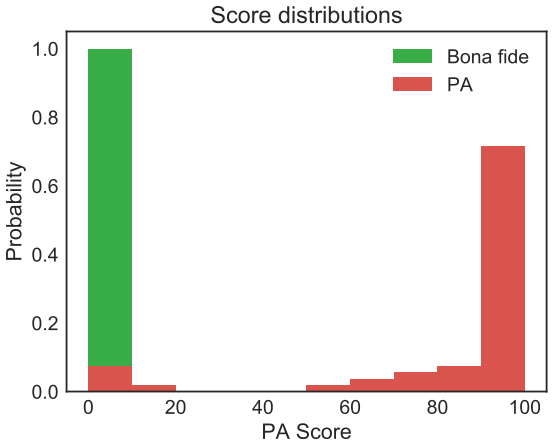

Figure 5: Normalised score distributions for bona fide (green) and PA (red) presentations. As it can be observed, only three scores yielded by PAs overlap with the scores stemming from bona fide presentations.

- overlays fabricated with dragon skin and urethane;
- fingerprints printed on regular matte paper and on transparency paper.

Regarding the bona fide samples, in order to maximise their variability, they have been captured from all fingers, from the little one to the thumb. For each bona fide and PAI, between two and three samples have been acquired.

Following common practices to achieve a balanced training of the classifier, the same number of samples (six) have been used for each class, namely: bona fides and PAs. For the latter, the following six different materials have been chosen: dragon skin overlay, urethane overlay, playdoh finger, printed fingerprint on matte paper, printed fingerprint on transparency paper and silly putty finger. Since in total twelve different materials have been analysed, the remaining six PAIs can be considered unknown attacks, thereby allowing us to test the robustness of the classifier in this more challenging scenario.

Finally, in compliance with the ISO/IEC IS 30107-3 on Biometric presentation attack detection - Part 3: Testing and Reporting [22], the following metrics are used to evaluate the performance of the PAD method:

- Attack Presentation Classification Error Rate (APCER): percentage of attack presentations wrongly classified as bona fide presentations.
- Bona Fide Presentation Classification Error Rate (BPCER): percentage of bona fide presentations wrongly classified as presentation attacks.

3.2 Results

The score distribution for the bona fide presentations (green) and the PA (red) are depicted in Fig. 5. As it may be observed, there is a very small overlap between both distributions: whereas all bona fide presentations achieve scores lower than 10 (i.e., less than 10% of their

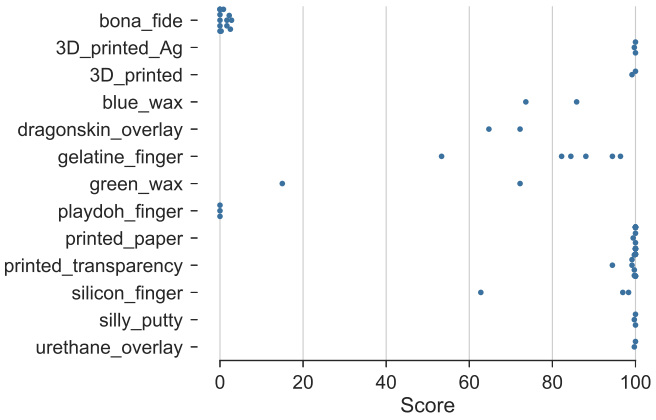

Figure 6: Scores yielded by each presenation, ordered by PAI species or bona fide along the y axis.

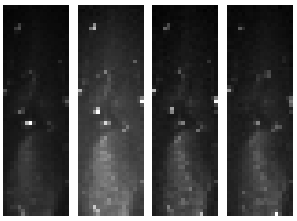

(a) (b) (c) (d)

Figure 7: Example of a noisy image for the green wax finger, captured at (a) 1200 nm, (b) 1300 nm, (c) 1450 nm, and (d) 1550 nm.

pixels are incorrectly classified as non-skin), only a single PAI (the playdoh finger), from which three samples have been acquired, achieves such a low score. Therefore, both classes are easily separable.

In order to further analyse the performance of the PAD method, all scores are presented in Fig. 6, classified along the y axis according to their origin: either bona fide or a particular PAI. For most of the bona fide samples, the score is 0 (i.e., all pixels are classified as skin), and all scores lie under 3 (i.e., less than 3% of the pixels are incorrectly classified).

Regarding the different PAIs considered, it may be observed that all playdoh finger samples achieve also a score of 0. This is due to the similarity in the responses to the selected wavelengths with respect to the bona fides, as it may be seen in Figs. 2b to 2e (bona fide) and Figs. 3i to 3j (playdoh PAI). On the other hand, all the remaining PAIs can be discriminated from bona fides with a threshold below 10. Moreover, most of the remaining PAIs yield scores over 80 (i.e., over 80% of the pixels are correctly classified as non-skin). Among the exceptions one of the green wax finger samples shows a remarkably lower score (15 vs. 75 – 85 for the remaining samples). This is due to a noisy acquisition, as shown in Fig. 7. In spite of that fact, the PA is still correctly detected.

In summary, the overall performance of the PAD algorithm for decision threshold on 5 (i.e.,

at least 5% of the pixels in the ROI are classified as non-skin) yields an APCER = 5.7% for BPCER = 0%, thereby granting a highly secure (low APCER) and convenient (low BPCER) system.

It should be finally highlighted that all unknown PAIs (i.e., both 3D printed fingerprints, blue and green wax fingers, gelatine finger and silicone finger) are correctly detected as PAs. That means the classifier is robust to all previously unseen attacks considered, thereby proving the soundness of the approach. Such success is due to the fact that most materials, excluding the orange playdoh tested, exhibit a similar behaviour to each other for the selected SWIR bands, and at the same time different from the bona fides. We may thus conclude that the proposed sensor and algorithm are able to detect unknown attacks.

4 Conclusions

We have presented a novel fingerprint presentation attack detection approach based on spectral signatures extracted from SWIR images. Contrary to other high performing PAD algorithms, since classification is performed in a pixelwise fashion, very few samples are required for training (six bona fides and six PAIs in our experimental evaluation). This allows for a robust training in spite of the lack of freely available databases comprising a wide variety of PAIs.

The experiments have shown that all bona fide samples are correctly classified as such, thus yielding a highly convenient system. On the other hand, all materials except for playdoh are also detected as PAs. Even unknown attacks are correctly detected, thereby showing the promising performance of the proposed method. This is due to the similarities exhibited by most PAIs at the selected SWIR bands, and their difference with respect to the bona fides behaviour.

In addition, given the reduced number of pixels of the final ROI considered (360 px, see Fig. 4), the low dimensionality of the feature vector (6) and the use of an SVM classifier, samples are classified in a fast and efficient manner. We can thus conclude that the proposed PAD method can be utilised for real-time applications.

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 to both known and unknown attacks. We will also explore other feature extraction techniques to deal with the undetected PAIs (playdoh finger) including both other algorithms trained on the same SWIR data and other acquisition technologies.

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.

References

[1] A. K. Jain, Technology: Biometric recognition, Nature 449 (207) 38–49.

[2] N. Ratha, J. Connell, R. Bolle, Enhancing security and privacy in biometrics-based authentication systems, IBM Systems Journal 40 (3) (2001) 614–634.

[3] ISO/IEC JTC1 SC37 Biometrics, ISO/IEC 30107-1. Information Technology - Biometric presentation attack detection - Part 1: Framework, International Organization for Standardization (2016).

[4] J. Galbally, M. Gomez-Barrero, Presentation attack detection in iris recognition, in: C. Busch, C. Rathgeb (Eds.), Iris and Periocular Biometrics, IET, 2017.

[5] J. Galbally, M. Gomez-Barrero, A review of iris anti-spoofing, in: Proc. Int. Workshop on Biometrics and Forensics (IWBF), 2016.

[6] E. Marasco, A. Ross, A survey on antispoofing schemes for fingerprint recognition systems, ACM Computing Surveys (CSUR) 47 (2) (2015) 28.

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

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

[9] S. Marcel, M. S. Nixon, S. Z. Li (Eds.), Handbook of Biometric Anti-Spoofing, Springer, 2014.

[10] TABULA RASA, Trusted biometrics under spoofing attacks (2010).
URL http://www.tabularasa-euproject.org/

[11] BEAT, Biometrics evaluation and testing (2012).
URL http://www.beat-eu.org/

[12] ODNI, IARPA, IARPA-BAA-16-04 (thor) (2016).
URL https://www.iarpa.gov/index.php/research-programs/odin/odin-baa

[13] Livdet - liveness detection competitions (2009–2017).
URL http://livdet.org/

[14] Y. Wang, X. Hao, Y. Hou, C. Guo, A new multispectral method for face liveness detection, in: Proc. Asian Conf. on Pattern Recognition (ACPR), 2013, pp. 922–926.

[15] H. Steiner, S. Sporrer, A. Kolb, N. Jung, Design of an active multispectral SWIR camera system for skin detection and face verification, Journal of Sensors 2016.

[16] R. K. Rowe, K. A. Nixon, P. W. Butler, Multispectral Fingerprint Image Acquisition, Springer London, 2008, pp. 3–23.

[17] S. Chang, K. Larin, Y. Mao, W. Almuhtadi, C. Flueraru, Fingerprint spoof detection by NIR optical analysis, in: State of the Art in Biometrics, InTech, 2011, pp. 57–84.

[18] C. Hengfoss, A. Kulcke, G. Mull, C. Edler, K. Püschel, E. Jopp, Dynamic liveness and forgeries detection of the finger surface on the basis of spectroscopy in the 400–1650 nm region, Forensic science international 212 (1-3) (2011) 61–68.

[19] J. A. Jacquez, J. Huss, W. McKeehan, J. M. Dimitroff, H. F. Kuppenheim, Spectral reflectance of human skin in the region 0.7–2.6 μ, Journal of Applied Physiology 8 (3) (1955) 297–299.

[20] H. Steiner, A. Kolb, N. Jung, Reliable face anti-spoofing using multispectral SWIR imaging, in: Proc. Int. Conf. on Biometrics (ICB), 2016, pp. 1–8.

[21] ISO/IEC JTC1 SC37 Biometrics, ISO/IEC DIS 30107-2. Information Technology - Biometric presentation attack detection - Part 2: Data formats, International Organization for Standardization (2017).

[22] ISO/IEC JTC1 SC37 Biometrics, ISO/IEC IS 30107-3. Information Technology - Biometric presentation attack detection - Part 3: Testing and Reporting, International Organization for Standardization (2017).