Mobile Face Recognition Systems: Exploring Presentation Attack Vulnerability and Usability

Heinz Hofbauer[1] • Luca Debiasi[1] • Andreas Uhl[1]

[1]Multimedia Signal Processing and Security Lab, University of Salzburg, Austria, {hofbauer, ldebiasi, uhl}@cs.sbg.ac.at

Abstract

We have evaluated face recognition software to be used with hand held devices (smartphones). While we can not go into specifics of the systems under test (due to NDAs), we can present the results of our evaluation of liveness detection (or presentation attack detection), matching performance, and success with different complexity levels of attacks. We will contrast the robustness against presentation attacks with the systems usability during regular use, and highlight where currently state of commercial of the shelf systems (COTS) stand in that regard. We will look at the results specifically under the tradeoff between acceptance, linked with usability, and security, which usually negatively impacts usability.

Contents

1 Introduction **2**

2 Related Work **2**

3 Usability **3**
3.1 Usability and Baseline 3
3.2 Usability and Baseline Outdoors 3
3.3 Discussion 4

4 Presentation Attack: Replay Attacks **4**
4.1 Discussion 5

5 Presentation Attack: Masks **6**
5.1 Discussion 6

6 Conclusion **6**

1 Introduction

We were tasked by a company with evaluating the usability and security of face recognition systems which work by recording a selfie (self-portrait) on a smartphone. The matching was done on the server side, but liveness detection was done on the smartphone. The company ran the servers, provided the hardware and software. The whole project was on a rather tight time schedule (due to license lease time), so we could only conduct a limited number of experiments with a limited number of people. Nonetheless, the results were rather interesting and we wanted to share them.

That said, this is not a very technical paper. It is more a recording of our experience with the software/devices/processes. The main incentive to share this information is to showcase certain problems which do not happen in a typical "lab setup". Shortcomings in algorithms or implementation can be detrimental to the adoption by industry or acceptance by users and it can occasionally lead to interesting research questions too. In this paper we will present our experiments and findings and comment on how research might help.

We fill focus more on the what is of interest to us as researchers and less on implementation details, except where the used protocol might impact the research side. That said, we would like to point out that software implementations, even only research software for reproducible research, should be built with corner cases in mind to allow for testing on more difficult test sets[1].

Limited Tests: Due to time constraints, only a short license lease time during which to test the systems was granted, we could only afford a very limited number of tests. Specifically, most test were only performed by a single user. The number of attempts was also rather low, usually 10 to 20 repeats per test. Yet, even with such a limited number of tests we could find counterevidence regarding the security of the systems.

The goal in all these tests is to have a method to unlock the device or otherwise verify the user of the device when the user cooperates. What is important to companies is that this process is secure on the one hand, but also fast and annoyance free for the user. If this latter part is not given, an adoption of the system by users is less likely.

As such we will look at the security of the two systems under test, PassiveSys and ActiveSys, with the goal of unlocking the device with minimal fuss on the part of the user.

The paper is structured as follows, Section 2 gives an overview of presentation attacks and their detection as it relates to the matter at hand. Section 3 establishes a baseline when genuine traits are presented to the systems under test. Section 4 will attack the test with replay type attacks and Section 5 will use more sophisticated replicas of the biometric traits to circumvent the system. Finally, Section 6 will summarize our findings and conclude the paper.

[1] While we will delve no deeper into this we just would like to note that we managed to crash the server because of the floral pattern design on a user's shirt worn during testing

2 Related Work

Smartphones are ubiquitous and so is the widespread adoption of biometric traits to unlock the device by verifying the identify of the user. In recent years, a certain trend from using fingerprints towards face detection can be observed. This trend has renewed the interest in attacks, and the prevention thereof, against such biometric systems.

A specific attack is the presentation, also known as direct or spoofing, attack. It can be separated into two categories [1]: (1) active imposter presentation attacks, where the attacker tries to claim a foreign identity; and (2) concealer presentation attacks, where an attacker tries to not be recognized by a system. Presentation attacks can be used against identification as well as verification modes.

Presentation attacks (PA) can also be differentiated by the source of the presentations attack instrument (PAI): (1) artificial, which is a non-human material sourced from humans, e.g., masks, printouts, images; (2) human traits, parts of dead bodies, modified faces, forced presentation by unconscious persons and so on.

To prevent such attacks, a presentation attack detection (PAD) system, also referred to as liveness detection, is employed. The primary focus of research is artificial presentation but, as is evident in the term liveness detection, overlaps with parts of the human trait PAI categorization.

There are different kinds of (face detection) PADs, some are hardware reliant while others are not, some use still images and others video. The number of different PAD methods is long, thus we will only give a brief list of methods without going into them too much: blink detection ([2, 3]), challenge response ([4–6]), texture based ([7–9]), dynamic texture based (video) ([10, 11]) or movement based ([12–14]). For more details, the reader is referred to the respective papers.

The target application of our tests was to unlock the device with the presented biometric trait (face). The operation mode, in terms of biometry, is always verification since the identify is implied (the owner of the cell phone). Presentation attacks also try to unlock the device and are consequently also done in verification mode. The presentation attack instrument is artificial only. While there are more types of PAIs, and a lot of further differentiation by subtype, we only gave related literature to the modes suspected to be employed in the devices we test. Specifically, ActiveSys certainly uses blink detection and challenge response methods. PassiveSys's modes are all passive, i.e., no cue based user interaction is required, using image, and we strongly suspect video, and thus has to rely on texture based image and video as well as movement features for PAD.

Please note: In the following sections we will present tables with results. These results are in the form of success rate of the liveness detection (LD) and the match rate (MR), which relate to the reporting as specified in ISO/IEC 30107-3 [15] as follows: In case the presented trait was genuine, the bona fide presentation classification error (BPCER) can be calculated as $\text{BPCER} := 1 - \text{LD}$. In case of presentation attacks the attack presentation classification error rate (APCER) can be calculated as

APCER := 1 − LD. Likewise, the false non match rate (FNMR) for genuine presentation is FNMR := 1 − MR and the impostor attack presentation match rate (IAPMR) for presentation attacks is IAPMR := MR.

3 Usability

For usability, we look at the basic modes provided by the software. With these modes we get a baseline for further tests and presentation attacks. We evaluated two software systems, denoted PassiveSys and ActiveSys, both have a separate step for detecting liveness and matching the probe and gallery image.

PassiveSys could operate with five different modes, which only impact liveness detection. No further detail on what is different was provided to us, but on-screen notes gave clues on what is required for the liveness detection. Modes are: *video*, unclear conditions but seems to take a video; *lessvid*, seems to be a less stringent version of *video*; *image*, simply takes a picture. One mode was not used because we could never pass liveness detection. Another mode was designed to use the rear camera and an operator to identify a second person, this was not used because the goal is to unlock the device (single user operation).

ActiveSys allows four liveness detection modes: None; *blink*, user has to keep still and blink on cue; *arrow*, requires turning the head to steer an arrow along a line to a target, when the arrow and target align the user has to blink; *blink+arrow*, a combination of both modes. We will not give separate results for None and the *blink+arrow* combination since the modes are simply executed one after the other.

3.1 Usability and Baseline

To get a baseline for the systems, we created two test sets, one where the gallery images is from a user with glasses and one where the user does not wear glasses.

The results are given in Table 1, split for system and liveness detection type. It can be seen that the presence of glasses in the image increases the error rate of the liveness detection. It is also interesting to see that the matching always worked when liveness detection was passed. However, even for probe images without glasses certain modes did reject a lot of attempts, *video* overall rejected almost 72% of all attempts. Also interesting is that *arrow* seems to reject less attempts than *blink*, even though the task is more complicated. The modes of PassiveSys on the other hand behave as expected, the more complicated method reject more attempts, i.e., video based reject more than image base liveness detection modes.

What also resulted from these experiments, which is not reflected in the table, is the insight that failure of longer modes, like the *arrow* or *blink+arrow* modes for ActiveSys which took several seconds per attempt, became frustrating very fast.

3.2 Usability and Baseline Outdoors

We suspected that the failure to detect images with glasses as alive was due to reflection of light on the glasses. The results in Table 1 were obtained from experiments in a well lit room. To further test the impact of light on the liveness detection and to expand the baseline to the outdoors, we performed another test in natural sunlight, during a bright day.

This experiment was conducted without glasses and the results are given in Table 2. The clear impact of lighting conditions on the liveness detection is quite drastic, *video* and *lessvid* failed to detect anything as alive and *image*, *blink* and *arrow* all had reduced number successful attempts. However, it should also be noted that the actual verification always worked when the liveness detection was passed. This might be a benefit of the aggressive screening during liveness detection, which is

Table 1: Baseline for the ActiveSys and PassiveSys. Results are split between liveness detection test (LD) and verification results (Match). The presence of glasses in the probe (Pr.) and gallery (Gal.) images is given as well.

(a) Baseline for PassiveSys for modes *video*, *lessvid*, *image*.

video

Pr.	Gal.	LD	Match
yes	yes	0/20	0/20
no	no	13/20	13/20
no	yes	10/20	10/20
yes	no	0/20	0/20

lessvid

Pr.	Gal.	LD	Match
yes	yes	4/20	4/20
no	no	18/20	18/20
no	yes	12/20	12/20
yes	no	9/20	9/20

image

Pr.	Gal.	LD	Match
yes	yes	6/20	6/20
no	no	20/20	20/20
no	yes	18/20	18/20
yes	no	19/20	19/20

(b) Baseline for ActiveSys for modes *blink* and *arrow*.

blink

Pr.	Gal.	LD	Match
yes	yes	20/20	20/20
no	no	16/20	16/20
no	yes	12/20	12/20
yes	no	12/20	12/20

arrow

Pr.	Gal.	LD	Match
yes	yes	18/20	18/20
no	no	20/20	20/20
no	yes	20/20	20/20
yes	no	20/20	20/20

Table 2: Performance during bright sunlight outdoors. Results are split between liveness detection test (LD) and verification results (Match). Facing was either towards the sun or away from the sun.

(a) PassiveSys split for modes.

video and *lessvid*

Facing	LD	Match
towards	0/20	0/20
away	0/20	0/20

image

Facing	LD	Match
towards	10/20	10/20
away	20/20	20/20

(b) ActiveSys split for modes.

blink

Facing	LD	Match
towards	14/20	14/20
away	9/20	9/20

arrow

Facing	LD	Match
towards	10/20	10/20
away	19/20	19/20

Table 3: Liveness Detection under studio light for different light positions (Dir.) and intensities. Light was diffused or undiffused as a spot light. Entries are the number of success based on 10 attempts per setting.

			LD under Intensities					
			spot			diffuse		
System	Mode	Dir.	1.0	3.0	6.0	1.0	3.0	6.0
PassiveSys	*video*	front	0	0	0	0	0	0
PassiveSys	*lessvid*	front	4	2	3	10	5	5
PassiveSys	*image*	front	8	8	9	10	9	10
ActiveSys	*blink*	front	8	8	6	2	5	3
ActiveSys	*arrow*	front	9	9	10	10	7	8
PassiveSys	*video*	side	0	0	0	0	0	0
PassiveSys	*lessvid*	side	3	4	2	7	5	0
PassiveSys	*image*	side	9	8	9	10	7	8
ActiveSys	*blink*	side	3	4	4	4	4	6
ActiveSys	*arrow*	side	3	5	1	10	5	3
PassiveSys	*video*	back	0	0	0	0	0	0
PassiveSys	*lessvid*	back	5	3	1	4	3	1
PassiveSys	*image*	back	5	9	6	2	0	1
ActiveSys	*blink*	back	6	7	5	7	5	3
ActiveSys	*arrow*	back	9	10	10	10	9	10

not necessarily a bad thing since early failure is less costly in terms of time to failure.

To get a more reproducible, and finer grained, version of the light test we set up a dimmed room with a studio light (Helios 300p) shining at the user from the front, side or back at a distance of roughly 1m. The light levels were adjustable and were set to 1, 3 and 6 (from a maximum setting of 6) and we investigated spot and diffuse (diffused with bleached 80g/m^2 paper) light to simulate a clear or cloudy day, one such experiment is depicted in Figure 1a. The results are given in Table 3, verification is not given separately since every time liveness was detected the user was also correctly verified. Results given are the successful unlock attempts from a set of 10 attempts per parameter set.

The results from the controlled tests show quite nicely the influence of light on the different modes. All modes are affected to some degree and for the most part how they are affected makes sense, e.g., higher effect the stronger the light is, spot light has a higher effect than diffuse light and so on. The one difference is the direction, frontal light illuminates the subject unlocking the device so has the least influence, but the higher reduction of side illumination over backlight is somewhat surprising. The sidelight usually results in a very uneven illumination, one face side in shadow the other illuminated. Backlight should mess up the exposure settings of the camera and leave the whole face in shadow. The expectation therefore would be that the backlight exhibits worse performance than sidelight, which is not backed by experimental results.

3.3 Discussion

Time to failure and repeats can heavily impact the user experience. The overall time taken to unlock has to be acceptable to the user. Failures do not matter so much, so if failure and retry is fast and painless then the resulting user experience can still be good. However, if a long process fails and has to be retried, user satisfaction quickly fades. On a related note users try to help the system by doing the "right" thing to speed up the process. We can use this by making explicit what is required rather than letting the user guess. The user is a willing participant and will try to help as much as possible to speed up the unlocking process.

Table 4: Result of a replay attack. The number of successful attacks out of 20 attempts is given.

		Successes with Replay		
System	Mode	print	screen	video
PassiveSys	*lessvid*	0	1	0
PassiveSys	*image*	12	17	20
ActiveSys	*blink*	—	—	0
ActiveSys	*arrow*	—	—	5

Light and the outdoors environment. The impact of directional light on the liveness detection system is quite drastic and will make many of the modes under test unfeasible in practice. And while the matching worked well for all cases it is not clear if this is due to the aggressive liveness screening or robust matching algorithms.

4 Presentation Attack: Replay Attacks

The next logical step to test the security of the system was to perform a replay attack. That is, record an image or video and present that to the device instead of the genuine face. In a perfect world the liveness detection should reject every attempt.

To reduce the amount of data to display in tables, the *video* mode will no longer be used. Given its problems of rejecting images with glasses and strong light, it will likely never be used in practice either.

For the simple replay attack, we used a printed version of the image, the image displayed on a computer screen and a short video also displayed on the screen. The latter was used since both modes from ActiveSys require at least blinking and a bit of interaction in the case of *arrow*, simulated by turning the smartphone. The setup for the static image replay attacks and test of degradation types (see below) is shown in Figure 1b.

The results from this test can be seen in Table 4, again only liveness detection is given since verification was always successful when liveness was detected.

It is interesting to compare these results to the lighting results in Table 2. The same stringency which allows the detection of replay attacks adversely affects the usability in environments with bright lights. Overall, the expected result is present, higher quality/effort reproductions have a higher success chance, i.e., video is better than screen is better than print. And again the *arrow* mode is easier to pass than the *blink* mode, even though more 'user' interaction is required.

(a) Evaluation of the impact of frontal studio light.

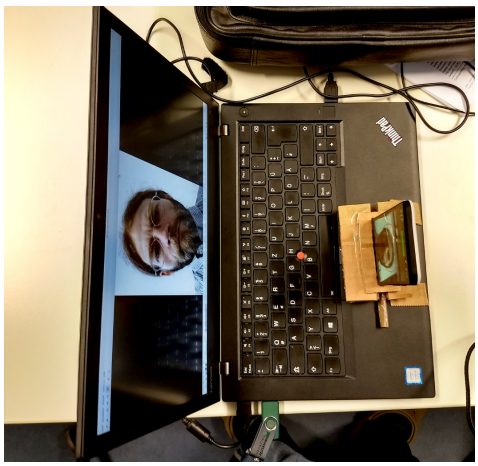

(b) Replay attacks for still images, same setup was also used for degraded images.

Figure 1: Different test setups.

Table 5: Result of a controlled degraded image replay attack. The number of successful attacks out of 10 attempts is given.

			Strength of Degradation		
System	Mode	Degradation	low	medium	high
PassiveSys	*image*	Noise	10	10	10
PassiveSys	*image*	Blur	10	10	6
PassiveSys	*image*	Resolution	10	10	6/0*

*liveness was detected 6 times, but verification was passed 0 times

Assuming that usability is a prime factor for a widespread adoption of such unlock systems, we will take a closer look at just how bad a recording still allows an unlock. From a practical perspective we will only look at the *image* mode. While *lessvid* would also be an interesting candidate, the mounting of such a replay attack is harder since a video has to be acquired, while *image* only requires a still image, i.e., a simple photograph. To simulate bad recording conditions we will add noise, blur the image and pixelate it to simulate a low resolution. The results are given in Table 5 for the *image* mode, Figure 2 illustrates the range of noise, blur and pixelation applied.

The clear result of these tests is that even a strongly degraded version of the image can penetrate the liveness detection of the *image* mode.

4.1 Discussion

An interesting tradeoff between usability and security can be observed in these experiments. Since usability is paramount for applicability, the security has the be reduced somewhat. However, this can be counteracted by user participation in activity assisted unlock modes like *arrow*. The drawback of such methods is that they take longer and require more attention from the user making a failure to unlock more annoying. This annoyance could hinder adoption of such schemes, which in turn would require a reduction in security and thus brings us full circle again. There is clearly a need for fast and reliable liveness detection methods.

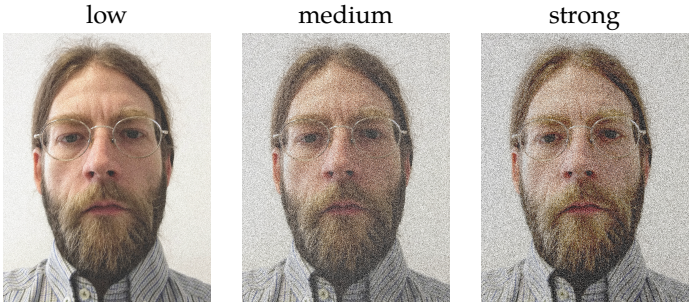

(a) RGB Noise

(b) Gaussian Blur

(c) Resolution

Figure 2: Illustration of Degradation types and strength for the replay attacks in Table 5.

Table 6: Presentation attack results for both mask types and the given modes and systems.

		Latex Mask		Resin Mask	
System	Mode	LD	M	LD	M
PassiveSys	*lessvid*	5	0	10	0
PassiveSys	*image*	10	0	20	0
ActiveSys	*blink*	0	0	10	0
ActiveSys	*arrow*	0	0	16	0

5 Presentation Attack: Masks

For these presentation attacks, we used a mask or mask-like presentation of the stolen biometric trait, created via photographs of the target's face. This was done to increase the chance of breaking interactive systems and give the impression of depth a 2D image might not convey.

We used two attack types: (1) a handcrafted 3D latex based mask by CREA FX[2]; and (2) a 3D-printed hard resin composite mask by ThatsMyFace[3]. Figure 3 show examples of the different masks.

Since the masks allow some interaction, we will again use both modes from ActiveSys as well as *image* and *lessvid* from PassiveSys. The results are given in Table 6 out of 20 attempts.

The obtained results are very interesting, especially in comparison to prior experiments. Where until now the liveness detection was relatively stringent and the verification always worked when the liveness detection was passed, the table has turned here. The liveness detection, which really should catch these cases fails and lets them pass, while the verification rejects the masks.

There is also quite the difference in mask quality, while the latex mask was handcrafted, took about three times as long to acquire and was four times as expensive as the 3d-printed resin mask, it performed worse.

5.1 Discussion

What is interesting here is that the relatively high cost only marginally increases the success rate. To illustrate this, let us have a look at the threat level model laid out in [16], briefly given in table 7.

Table 8 compares this to the results from our test, where success rate is the percentage of presentation which passed the liveness detection and verification. Usability is the chance of unlock by a genuine user under different conditions. What is most interesting is that Level C attacks, which are much more expensive and have a much higher preparation time, do not improve in success rate over Level B and Level A attacks.

From this table it also becomes clear that the PassiveSys system is basically unusable, either the usability of a given mode is low (*lessvid*) or the success rate of attack is high (*image*). The ActiveSys system is far better designed in this regard. A trade-off between reduced usability and higher security (*blink*) and higher usability at the cost of a potential Level B attack (*arrow*) can be observed.

[2] https://www.creafx.com/en/

[3] http://thatsmyface.com/custom-wearable-masks/

Table 7: Spoof presentation attacks separated by levels based on time, expertise, and equipment.

Threat	Level A	Level B	Level C
Time	short	>3 days	>10 days
Expertise	anyone	practice needed	extensive skill required
Equipment	readily available	requires planning	specialized
Biometric source	readily available	difficult to obtain	difficult to obtain
Example	paper print of image	paper mask or video	3D face reconstruction

Table 8: Comparison of threat level and success rate per mode and system. Usability, the chance of unlock by a genuine user is a combination of results from Table 1 and 2.

System	Mode	Attack	Threat	Success Rate	Usability
PassiveSys	*lessvid*	Image	Level A	5%	
PassiveSys	*lessvid*	Video	Level B	0%	44.2%
PassiveSys	*lessvid*	Mask	Level C	0%	
PassiveSys	*image*	Image	Level A	85%	
PassiveSys	*image*	Video	Level B	100%	77.5%
PassiveSys	*image*	Mask	Level C	0%	
ActiveSys	*blink*	Video	Level B	0%	69.2%
ActiveSys	*blink*	Mask	Level C	0%	
ActiveSys	*arrow*	Video	Level B	25%	89.2%
ActiveSys	*arrow*	Mask	Level C	0%	

6 Conclusion

What we have seen is that biometric verification for the systems under test seems to work well. However, it is unclear if this is in part due to the strict liveness detection. While this may seem an odd differentiation, we have also seen that a strict liveness detection can reduce usability. At times this reduction can be quite drastic and based on plain and simple factors, like wearing glasses or trying to unlock the device during a bright day. As such, a step towards a higher usability and consequently user satisfaction and acceptance, would be to tweak the liveness detection to be less strict in such cases. However, if this has a negative effect on the matching performance, nothing is gained in terms of usability at the cost of security.

That said, the liveness detection of both tested systems does a relatively good job of screening attacks. Again, this success in screening attacks is at the cost of usability. While this trade-off is fine in theory, the practical impact is quite high,i.e., PassiveSys reduced the chance of success for genuine presentations to less than 50% and could still be successfully attacked.

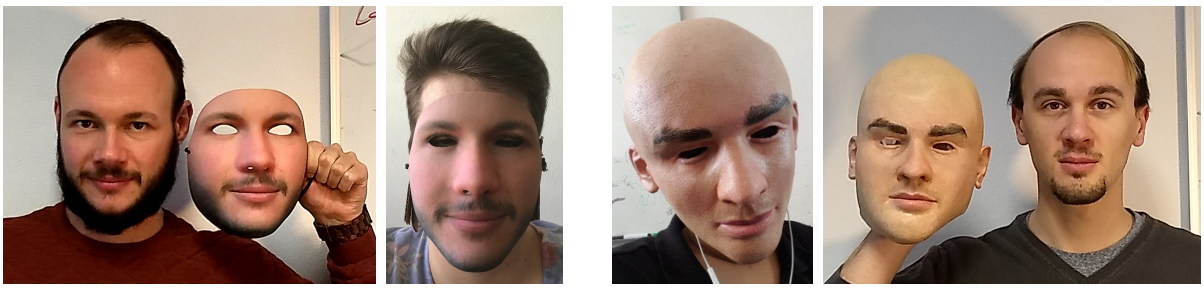

Resin composite Latex

Figure 3: Masks used for presentation attacks. The sources of the biometric traits hold their replicas and imposter wearing the replicas during an attack attempt.

While ActiveSys fared better, it also had to reduce the usability to around 70% to prevent attacks. There is clearly ample room for improvement.

Regarding the attacks, it was interesting to see that the most expensive and time consuming attacks, specially created facial masks, fared worse than relatively simple printed image or video presentation attacks.

Acknowledgements

This work has been partially supported by Veridos Inc.

References

[1] R. Raghavendra and C. Busch, "Presentation attack detection methods for face recognition systems: A comprehensive survey," *ACM Computing Surveys*, vol. 50, no. 1, 2017 (cit. on p. 2).

[2] P. Gang, S. Lin, W. Zhaohui, and L. Shihong, "Eyeblink-based anti-spoofing in face recognition from a generic webcamera," ser. Proceedings for IEEE 11th International Conference on Computer Visio, 2007 (cit. on p. 2).

[3] M. Chrzan, "Liveness detection for face recognition," Master's thesis, Masaryk University, Factory of Informatics, 2014 (cit. on p. 2).

[4] K. Kollreider, H. Fronthaler, and J. Bigun, "Verifying liveness by multiple experts in face biometrics," ser. IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, 2008 (cit. on p. 2).

[5] A. Ali, F. Deravi, and S. Hoque, "Directional sensitivity of gaze-collinearity features in liveness detection," in *4th International Conference on Emerging Security Technologies*, 2013 (cit. on p. 2).

[6] D. Smith, A. Wiliem, and B. Lovell., "Face recognition on consumer devices: Reflections on replay attacks," *IEEE Transactions on Information Forensics and Security*, vol. 10, no. 4, 2015 (cit. on p. 2).

[7] N. Kose and J. Dugelay, "Classification of captured and recaptured images to detect photograph spoofing," in *International Conference on Informatics, Electronics Vision (ICIEV'12)*, 2012 (cit. on p. 2).

[8] J. Yang, Z. Lei, S. Liao, and S. Z. Li, "Face liveness detection with component dependent descriptor," in *International Conference on Biometrics (ICB'13)*, 2013 (cit. on p. 2).

[9] Z. Boulkenafet, J. Komulainen, and A. Hadid, "Face spoofing detection using colour texture analysis," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 8, 2016 (cit. on p. 2).

[10] I. Chingovska, A. André, and S. Marcel, "On the effectiveness of local binary patterns in face anti-spoofing," in *Proceedings of the International Conference of the Biometrics Special Interest Group (BIOSIG '12)*, 2012 (cit. on p. 2).

[11] A. A. Tiago de Freitas Pereira, J. M. D. Martino, and S. Marcel, "LBP-TOP based countermeasure against face spoofing attacks," in *Computer Vision - ACCV 2012 Workshops*, ser. Lecture Notes in Computer Science, vol. 7728, 2012 (cit. on p. 2).

[12] M. D. Marsico, M. Nappi, D. Riccio, and J. Dugelay, "Moving face spoofing detection via 3d projective invariants," in *5th IAPR International Conference on Biometrics (ICB'12)*, 2012 (cit. on p. 2).

[13] A. Anjos, M. M. Chakka, and S. Marcel, "Motion-based countermeasures to photo attacks in face recognition," *IET Biometrics*, 3 2014 (cit. on p. 2).

[14] A. Pinto, W. Schwartz, H. Pedrini, and A. Rocha, "Using visual rhythms for detecting video-based facial spoof attacks," *EEE Transactions on Information Forensics and Security*, vol. 10, no. 5, 2015 (cit. on p. 2).

[15] ISO/IEC 30107-3, *Information technology - biometric presentation attack detection - part 3: Testing and reporting*, 2017 (cit. on p. 2).

[16] S. Schuckers, "Presentations and attacks, and spoofs, oh my," *Image and Vision Computing*, vol. 55, 2016, Recognizing future hot topics and hard problems in biometrics research. DOI: 10.1016/j.imavis.2016.03.016 (cit. on p. 6).