Spoofing 2D Face Recognition Systems with 3D Masks

Nesli Erdogmus, Sébastien Marcel
Idiap Research Institute
Centre du Parc - rue Marconi 19 CH-1920 Martigny, Suisse
{nesli.erdogmus,sebastien.marcel}@idiap.ch

Abstract:
Vulnerability to spoofing attacks is a serious drawback for many biometric systems. Among all biometric traits, face is the one that is exposed to the most serious threat, since it is exceptionally easy to access. The limited work on fraud detection capabilities for face mainly shapes around 2D attacks forged by displaying printed photos or replaying recorded videos on mobile devices. A significant portion of this work is based on the flatness of the facial surface in front of the sensor. In this study, we complicate the spoofing problem further by introducing the 3rd dimension and examine possible 3D attack instruments. A small database is constructed with six different types of 3D facial masks and experimented on to determine the right direction to study 3D attacks. Spoofing performance for each type of mask is assessed and analysed thoroughly using two Gabor-wavelet-based algorithms.

1 Introduction

Automatic recognition of human biometric traits has numerous important advantages over conventional methods like passwords or ID cards [JR08] and hence, it has become a vast research field today as the need and investment for access control systems grow continuously. Among these traits, face recognition stands out with its favourable reconciliation between convenience and reliability. Unfortunately it suffers security vulnerabilities because it is easy to access to face samples and subsequently, to devise spoofing attacks.

Spoofing attack is the act of presenting a fake biometric evidence to a system in order to achieve authentication [NAR08]. Forging such an attack is relatively simple for facial recognition systems since the photographs or videos of valid users can be acquired from internet or captured at a distance. Attackers can attempt to penetrate by displaying printed photos or replaying recorded videos on mobile devices in front of the sensors. Since these are the most common, easiest and cheapest methods to circumvent face recognition systems [BAF+12], the counter measure studies primarily shape around them. Recently, several papers that enable comparison of various counter attack algorithms are published, providing public databases and reproducible works [CAM+11, AM11, CAM12].

Many anti-spoofing techniques analyse the texture of the captured image, mainly based on the assumption of printing artifacts [BNGS10] and/or blurring [LWTJ04] which rely on the printed image or display quality.

Another group of methods that appears as a separate or complementary measure, tries

to detect liveness of the face based on live-face specific movements such as eye blinking [PSWL07]. However, this kind of methods will definitely fail in the case of video replay attacks or even more simply, photographic masks which are high resolution photographs worn on face with eyes and mouth regions cut out, as illustrated in [KFB08].

Finally, there exist motion analysis techniques to detect spoofing attacks based on the fact that planar objects like papers or screens move in a significantly different way than the real faces. For instance, in [KFB05], the trajectories of single parts of faces are analysed to distinguish between live and spoofed ones. In a similar manner, Marsico et al. [DMNRD12] exploit the facial shape and detect attacks by computing geometric invariants of a set of automatically located facial points. The 3D nature of the captured face can also be perceived by employing additional devices. Even though they were mostly considered to be expensive, this view is bound to grow obsolete with the introduction of affordable consumer depth cameras. In [TDM07], 3D data acquired with a low-cost sensor is utilized to localize face and at the same time to test the "face-ness", rendering the system invulnerable to spoofing attacks. This claim is pretty valid, since differentiating a real face from a planar surface with a depth sensor is quite straightforward.

On the other hand, the advancements in depth sensing technologies are counteracted by similar progress trends in 3D facial mask manufacture. Taking the face spoofing attacks one step, namely one dimension, further, 3D masks introduce new challenges for counter measure studies. As discussed above, the majority of the counter-measures developed for 2D spoofing fails to function properly in the case of masks.

In [ZYL+12] Zhang et al. states that since usually it is too expensive to produce client-like masks, massive usage of masks rarely appears in the literature. To the best of our knowledge, there have been three papers published in this field, which aim to detect 3D mask attacks using multi-spectral lighting [KNYY09, ZDLL11]. In both approaches, the authors try to classify human skin and non-skin by using Lambertian model to analyse multi-spectral reflectance properties of face. In [KNYY09], reflectance of human skin and different mask materials(silicon,latex and skin-jell) is measured on the forehead region for visible and infra-red light at different wavelengths. A 2D feature vector is proposed under 850 and 685 nm illumination which reaches a classification accuracy of 96.77%. In a similar manner, Zhang et al. [ZDLL11] select two most discriminative wavelengths and train an SVM classifier to learn real face and mask distributions at multi-distances. An average detection accuracy of 89.18% is achieved. Lastly, in a recent publication [KD13], the authors propose an LBP-based counter-measure against 3D printed masks. The spoofing detection accuracies are reported to be 88% and 86% for color and depth images. There are two main shortcomings in those studies: As later stated by the authors of [ZDLL11] in [YZL+12], the first limitation of the first two techniques is that they are not very convenient due to their special expensive hardware requirements. Secondly and most importantly, no analysis on the spoofing performances of the 3D masks is included.

In this paper, our purpose is to fill this gap by shedding some light on how high the spoofing performances can get for different types of masks. We believe that developing counter-measures for an impractical attack would be a waste of resources. Our experiments show that silicone masks whose reflectance properties are analysed in depth [KNYY09, ZDLL11] and 3D printed masks which are examined in [KD13] to develop anti-spoofing

measures, in fact, can not deceive 2D face recognition systems that well. It is important to emphasize that our study merely focuses on examination of spoofing performances and counter measure design is out of its scope. Despite its small size, the collected dataset offers crucial preliminary findings which a larger-scale research can be based on.

2 3D facial masks

The production of 3D facial masks spans a wide variety of materials and methods. It can range from an easy/low-cost way of printing a 2D photo on a deformable surface like cloth to a more difficult/expensive method of obtaining and printing a 3D model of a person.

Printing photo on a cloth is not a very promising technique to spoof a 2D face recognition system since the image will be distorted when applied to the attackers face. In our studies, we analyse a more sophisticated way of creating 3D masks out of printed 2D face images. But first, we like to talk about a recent service called ThatsMyFace.com which specializes in facial reconstruction and in transforming 2D portraiture into 3D sculptures. Using this service one only has to upload 2D mugshots (one frontal and one profile) of a person and he can receive this person's constructed 3D face in his mailbox, either in a nice picture frame as a present, or on an action figure as a toy, or as a wearable life-size mask that can be easily used in a spoofing attack.

In our research, we include four types of masks from ThatsMyFace.com. The first type (M1) comes as a PDF paper-craft layout of the face in 3D and can be printed, cut, folded and glued together into a 3D face mask. The other three types of masks are made out of a hard resin composite in full 24-bit colour with a matte varnish. One is $\frac{1}{2}$ life-size (M2) while the other two are of real size (M3 and M4). In addition to these masks, we analyze two more mask types. The first one is printed from a real 3D scan of a person (M5). For the time being, the acquisition of a detailed 3D face model of a person without his permission is not very feasible since the scanning process mostly requires subject cooperation. But with the pace of advancements in 3D sensing technologies, it is still beneficial to analyse this kind of masks. The last type of mask is made of silicone that is produced using face mould (M6) and painted realistically.

3 Experiments and results

3.1 Face recognition systems

In order to test the spoofing performances of 3D face masks, an open source framework for standardized comparisons of face recognition algorithms is utilized [GWM12]. Among a variety of implemented methods in the framework, two gabor-wavelet-based algorithms are selected which do not require any training of a prior model, naturally avoiding the generalizability problem: Local Gabor binary pattern histogram sequences

	M1	M2	M3	M4	M5	M6	PA	Real
S1	7	-	19	-	-	-	20	12
S2	12	32	-	25	-	-	20	17
S3	-	-	-	-	-	60	22	9
S4	-	-	-	-	16	-	21	25

Table 1: Number of samples in the utilized database for each subject (S1,S2,S3,S4) and attack type (masks: M1,M2,M3,M4,M5,M6 and photo attack: PA) together with example face images from the database: The first row is composed of real faces. Samples of different life-size masks are given in the second row. The first two are ordered from ThatsMyFace.com with and without holes at the eyes and the nostrils, and the third is directly printed from a real 3D face model. Finally in the last row, paper-cut, half-size (from ThatsMyFace.com) and silicon masks produced using face mould are presented from left to right.

(LGBPHS) [ZSG+05] and Gabor graphs [WFKvdM97] with a Gabor phase based similarity measure [GHW12]. Additionally, due to the robustness of Gabor features against local distortions caused by variance of illumination, expression and pose, they have been successfully and extensively applied for face recognition.

For both systems, the faces are firstly geometrically normalized using manually labelled eye positions and histogram equalization is applied. For the LGBPHS algorithm, facial images are convoluted with a set of 40 Gabor wavelets and Local Binary Pattern histograms [AHP04] are calculated for non-overlapping 8×8 pixel blocks. Finally, 64 resulting histograms are concatenated into a final feature vector and compared using the χ^2 metric. In the second method, the Gabor jets collected in grid graphs are compared using both the magnitude and phase of the Gabor wavelet response. For more details, interested readers can also refer to open-source implementations of these methods in BOB [AESW+12].

3.2 Database

For this study, we constructed a small database that consists of 4 subjects. Apart from 4 gallery samples, 59 face images are collected for real attempts under controlled illumination conditions. Low similarity scores can be observed for genuine and spoofing samples under adverse conditions, clouding the judgement on the exact impact of the masks. This behaviour is undesirable when assessing the spoofing performances.

In order to perform the attacks, we obtained different types of masks for each subject. There are no types of mask which were manufactured for more than one person, except M1. For comparison reasons, we also included photo attack samples (PA). The number of samples for each type of attack and real and spoofing samples for all subjects with each type of attack are given in Table 1.

3.3 Spoofing performances

Before evaluating the 3D mask attacks, two face recognition systems are analysed and the operating point for the score threshold is decided to be at the equal error rate (EER) where false accept and false reject rates are equal[1]. With these settings, at the score threshold of 0.59, Gabor graphs algorithm achieves 1.64% EER. Falling a little behind, LGBPHS algorithm reaches 5.19% at the threshold of 1.69. The photo attacks at these thresholds achieve 78.31% and 97.59% success for Gabor graphs and LGBPHS algorithms, respectively.

The similarity scores for masks are analysed separately. For M1, the score distributions for genuine, impostor and attacker scores are obtained and analyzed. At the EER threshold, this type of masks reach 5.26% attack success rate. This is expectedly far behind the photo attacks. The paper-craft masks have distorted shapes and noisy textures. The edges of the patches are clearly visible.

If we break down the results even further and analyse the two M1 masks separately, it is observed that the mask for subject 2 performs much better (14.29%) and in fact none of the samples from subject 2 manage to penetrate. This may be due to several reasons: The first factor is illumination conditions. The planar facets on facial surface lead to specular reflections and create susceptibility to lighting. M1 for subject 2 is sampled under more ”ideal” conditions. The second factor is the process of production for these masks. The 3D shape for the masks are calculated from the 2D photos of the clients. Additionally, crafting them requires manual skills. Hence, the accuracy of the outcome may not be the same for all cases. Still, it does not change the fact that the chances are very low for these type of masks to spoof a good 2D face recognition system.

Scale invariance in 2D face recognition gives a remarkable advantage to M2 masks. Thanks to their compactness, attackers can easily carry and use them with a high chance of success and without drawing too much attention. At EER threshold these masks reach 78.12% success rate in both recognition systems, drawing very close to photo attacks.

Oddly, the realistic life-size mask M3 do not meet the expectations and perform very poorly. It reaches spoofing success rates of 21.05% and 15.79% for Gabor graph and LGBPHS systems, respectively. As is the case with M1, these masks highly depend on the input 2D images that are utilized to reconstruct the 3D shape and the quality and correctness of the final product.

M4, on the other hand, is the most successful attack among all types, reaching 100% success rate for both systems. In fact, the holes at the eyes are expected to be detrimental, since the attackers are not always able to align their eyes correctly. But, this argument is much strongly disproved by the obtained results. Instead, they made it clear that the spoofing success for 3D life-size masks is highly contingent on degree of accuracy in mask manufacture. In addition, again there is a substantial difference between the lighting conditions of M3 and M4. M4 is photographed in better conditions with respect to M3 and this may also have impact on the performances.

M5 is different from the previous masks in the sense that it has the exact shape of the client

[1]Longer version of this paper with additional figures and illustrations can be found at http://publications.idiap.ch/index.php/publications/show/2658

face, because it is directly printed from a real face scan using a 3D printer. In the case of a 3D face recognition system, it would be expected to achieve high attack success rates. However, in 2D face recognition, the texture is of higher importance. The printed mask is the first face model in the training set of the FRGC database [PFS+05]. The texture on the printed mask did not turn out to be very realistic. Still, the attack success rates are better than M3; 25.00% and 31.25% for Gabor graphs and LGBPHS algorithms, respectively.

Among all masks, M6 is the only one that has been studied to develop counter measures in the literature [KNYY09, ZDLL11]. Unfortunately, it is also the least successful. None of the spoofing attempts with the silicone mask manages to break into the two authentication systems. In score distributions, the silicon mask is observed to behave similarly to impostor (zero-effort) trials. The silicone masks may be used for evasion where the purpose is to hide ones identity. But this does not hold for spoofing, because in order to have a silicone mask of a valid user, the attacker needs to know the 3D shape of the user's face. Even if he did so, it requires painting skills to create a realistic texture on the mask surface. Briefly, for the time being, it is highly impractical to forge a 3D silicone mask attack and gain illegitimate access through 2D face recognition systems.

3.4 Conclusion

Even though the appearance of affordable customer depth cameras in the market facilitates face-ness detection and anti-spoofing for photo and video replay attacks, this advantage is counteracted with strong advancement trends in facial mask manufacture. As the thriving 3D mask production industry brings in new challenges for counter measure studies, the emphasis on this issue is still very limited. Two existing works [KNYY09, ZDLL11] try to differentiate human skin and silicon mask using spectral lighting, however neither of the papers contains an analysis on spoofing performances of the masks under discussion. The scarcity of works in this area is mainly due to the difficulty and costliness of constructing an extensive 3D mask database. In this paper, we present a preliminary study on different types of masks available in the market, paving the way for such a database and eventually more detailed exploration of the subject.

According the obtained results, spoofing attacks with silicone masks for which counter measures were proposed previously in [KNYY09, ZDLL11, KD13] are found to be the most ineffective and impractical. This outcome is solely enough to justify our purpose in this analytic study. The most successful attacks are performed by real life-size masks from ThatsMyFace.com with holes at the eyes and nostril (M4). Reaching 100% attack success rate, they are proven to impose the highest threat on 2D face recognition systems.

Based on these findings, as future work, we plan to collect a database of 3D spoofing attacks using masks of type M4. Further investigation is needed using a larger database with higher number of subjects and more real/fake access attempts. Additionally, in longer term, we aim to develop counter measures for this type of attacks. It would also be interesting to assess how these masks perform in front of depth sensors and investigate if it is possible to make use of the shape data in anti-spoofing. For this purpose, we aim to

employ both 2D and 3D devices in creating our 3D spoofing attack database.

Acknowledgements

The authors would like to express their thanks to the FP7 European TABULA RASA Project (257289) for its financial support and to Ralph Breithaupt for his collaboration.

References

[AESW+12] A. Anjos, L. El-Shafey, R. Wallace, M. Günther, C. McCool, and S. Marcel. Bob: a free signal processing and machine learning toolbox for researchers. In *ACM International Conference on Multimedia*, pages 1449–1452, 2012.

[AHP04] T. Ahonen, A. Hadid, and M. Pietikäinen. Face Recognition with Local Binary Patterns. In *European Conference on Computer Vision*, volume 3021 of *Lecture Notes in Computer Science*, pages 469–481. 2004.

[AM11] A. Anjos and S. Marcel. Counter-measures to photo attacks in face recognition: A public database and a baseline. In *International Joint Conference on Biometrics*, pages 1 –7, October 2011.

[BAF+12] B. Biggio, Z. Akhtar, G. Fumera, G.L. Marcialis, and F. Roli. Security evaluation of biometric authentication systems under real spoofing attacks. *IET Biometrics*, 1:11–24(13), March 2012.

[BNGS10] J. Bai, T.-T. Ng, X. Gao, and Y.-Q. Shi. Is physics-based liveness detection truly possible with a single image? In *IEEE International Symposium on Circuits and Systems*, pages 3425 –3428, June 2010.

[CAM+11] M.M. Chakka, A. Anjos, S. Marcel, R. Tronci, D. Muntoni, G. Fadda, M. Pili, N. Sirena, G. Murgia, M. Ristori, F. Roli, Junjie Yan, Dong Yi, Zhen Lei, Zhiwei Zhang, S.Z. Li, W.R. Schwartz, A. Rocha, H. Pedrini, J. Lorenzo-Navarro, M. Castrillon-Santana, J. Maatta, A. Hadid, and M. Pietikainen. Competition on counter measures to 2-D facial spoofing attacks. In *International Joint Conference on Biometrics*, pages 1 –6, October 2011.

[CAM12] I. Chingovska, A. Anjos, and S. Marcel. On the effectiveness of local binary patterns in face anti-spoofing. In *International Conference of Biometrics Special Interest Group*, pages 1 –7, September 2012.

[DMNRD12] M. De Marsico, M. Nappi, D. Riccio, and J. Dugelay. Moving face spoofing detection via 3D projective invariants. In *International Conference on Biometrics*, pages 73 –78, April 2012.

[GHW12] M. Günther, D. Haufe, and R. P. Würtz. Face recognition with disparity corrected gabor phase differences. In *International Conference on Artificial Neural Networks and Machine Learning*, volume 1, pages 411–418, 2012.

[GWM12] M. Günther, R. Wallace, and S. Marcel. An Open Source Framework for Standardized Comparisons of Face Recognition Algorithms. In *European Conference on Computer Vision Workshops and Demonstrations*, pages 547–556, 2012.

[JR08] AnilK. Jain and Arun Ross. Introduction to Biometrics. In A. K. Jain, P. Flynn, and A. Ross, editors, *Handbook of Biometrics*, pages 1–22. Springer US, 2008.

[KD13] N. Kose and J.-L. Dugelay. Countermeasure for the protection of face recognition systems against mask attacks. In *IEEE International Conference on Automatic Face and Gesture Recognition*, April 2013.

[KFB05] K. Kollreider, H. Fronthaler, and J. Bigun. Evaluating liveness by face images and the structure tensor. In *IEEE Workshop on Automatic Identification Advanced Technologies*, pages 75 – 80, October 2005.

[KFB08] K. Kollreider, H. Fronthaler, and J. Bigun. Verifying liveness by multiple experts in face biometrics. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 1 –6, June 2008.

[KNYY09] Y. Kim, J. Na, S. Yoon, and J. Yi. Masked fake face detection using radiance measurements. *Journal of the Optical Society of America A*, 26(4):760–766, 2009.

[LWTJ04] J. Li, Y. Wang, T. Tan, and A. K. Jain. Live face detection based on the analysis of Fourier spectra. In *SPIE 5404, Biometric Technology for Human Identification*, pages 296–303, 2004.

[NAR08] KristinAdair Nixon, Valerio Aimale, and RobertK. Rowe. Spoof Detection Schemes. In A. K. Jain, P. Flynn, and A. Ross, editors, *Handbook of Biometrics*, pages 403–423. Springer US, 2008.

[PFS+05] P. J. Phillips, P.J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview of the Face Recognition Grand Challenge. In *Computer Vision and Pattern Recognition Conference*, pages 947–954, 2005.

[PSWL07] G. Pan, L. Sun, Z. Wu, and S. Lao. Eyeblink-based Anti-Spoofing in Face Recognition from a Generic Webcamera. In *IEEE International Conference on Computer Vision*, pages 1 –8, October 2007.

[TDM07] F. Tsalakanidou, C. Dimitriadis, and S. Malassiotis. A Secure and Privacy Friendly 2D+3D Face Authentication System Robust Under Pose and Illumation Variation. In *International Workshop on Image Analysis for Multimedia Interactive Services*, page 40, June 2007.

[WFKvdM97] L. Wiskott, J.M. Fellous, N. Kuiger, and C. von der Malsburg. Face recognition by elastic bunch graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7):775–779, 1997.

[YZL+12] J. Yan, Z. Zhang, Z. Lei, D. Yi, and S.Z. Li. Face Liveness Detection by Exploring Multiple Scenic Clues. In *International Conference on Control, Automation, Robotics and Vision*, 2012.

[ZDLL11] Z. Zhang, D.Yi, Z. Lei, and S.Z. Li. Face liveness detection by learning multispectral reflectance distributions. In *IEEE International Conference on Automatic Face Gesture Recognition and Workshops*, pages 436 –441, March 2011.

[ZSG+05] W. Zhang, S. Shan, W. Gao, X. Chen, and H. Zhang. Local gabor binary pattern histogram sequence (lgbphs): A novel non-statistical model for face representation and recognition. In *IEEE International Conference on Computer Vision*, volume 1, pages 786–791, 2005.

[ZYL+12] Z. Zhang, J. Yan, S. Liu, Z. Lei, D. Yi, and S.Z. Li. A face antispoofing database with diverse attacks. In *International Conference on Biometrics*, pages 26–31, 2012.