

THE SWAX BENCHMARK: ATTACKING BIOMETRIC SYSTEMS WITH WAX FIGURES

Rafael Henrique Vareto, Araceli Marcia Saldanha, William Robson Schwartz

Smart Sense Laboratory, Department of Computer Science
Universidade Federal de Minas Gerais, Belo Horizonte, Brazil

ABSTRACT

A face spoofing attack occurs when an intruder attempts to impersonate someone who carries a gainful authentication clearance. It is a trending topic due to the increasing demand for biometric authentication on mobile devices, high-security areas, among others. This work introduces a new database named Sense Wax Attack dataset (SWAX), comprised of real human and wax figure images and videos that endorse the problem of face spoofing detection. The dataset consists of more than 1800 face images and 110 videos of 55 people/waxworks, arranged in training, validation and test sets with a large range in expression, illumination and pose variations. Experiments performed with baseline methods show that despite the progress in recent years, advanced spoofing methods are still vulnerable to high-quality violation attempts.

Index Terms— spoofing, presentation attack, face, wax figures

1. INTRODUCTION

The term *spoofing*, also referred to as *presentation attack*, represents a legitimate threat for biometric systems. The attack occurs when a malefactor attempts to pass him/herself off as someone who carries an advantageous authentication clearance. Intruders regularly employ falsified data to bypass security procedures and gain illegitimate access. As a countermeasure, new datasets are regularly released in the literature as an attempt to address the latest attack nuances and leverage upcoming researches [1, 2, 3].

The human face may have become the “universal” authentication biometric for several reasons, such as the increasing distribution of personal pictures on social networks, spread of surveillance cameras, and convenience to name a few [4]. Most spoofing attacks comprise facial pictures from users enrolled in recognition systems since face images are easily acquired due to their broad availability. A case in point occurred in China involving scammers posing with wax figures of prominent executives to lure in around 600K investors and embezzle almost US\$500 million [5]. The low-cost access to face images contributes to the increase of criminals designing attacks to be validated as authentic users, turning face spoofing into a popular way of deceiving people and biometric applications.

This work presents a database comprised of both real human and wax figure image/video in the interest of endorsing the problem of face spoofing detection. The proposed database, entitled Sense Wax Attack dataset (SWAX), is designed to investigate the problem in which a face media is presented to a system that must determine whether it categorizes a *bona fide* (authentic) or a counterfeit (attack) sample. The goal is not only to deliver a novel spoofing benchmark by specifying dataset protocols and usage requirements, but also enable the development of more robust anti-spoofing systems with the anticipation of unforeseen wax-based portrait attacks.



Fig. 1: A Korean celebrity on the left and a realistic sculpted wax dummy.

Although several datasets focus on spoofing detection, most engage in recapturing authentic images or videos in distinct mediums [6, 7, 8] by varying input sensors, attack types and capture conditions. Only recently have few researchers turned themselves to modeling emerging attack strategies [9, 10, 11], e.g. silicon masks, which suggests concealing a suspect’s real appearance or impersonating someone else’s identity. To the best of our knowledge, SWAX is the only dataset comprising both real and wax-modeled persons up to the present time. It provides means of studying attacks in unconstrained environments as it encompasses significant divergence in pose, expression, lighting, scene, and camera settings.

The proposed SWAX dataset consists of genuine and counterfeit samples for all available subjects. As illustrated in Figure 1, the database contains labeled photographs of characters, celebrities and public figures, chosen based on a list of waxworks obtained from a famous chain of wax museums. More precisely, this work aims at investigating face spoofing detection in realistic scenarios, whither there is little control over the images acquisition. In contemplation of fair algorithm comparisons, we provide four protocols for developing and evaluating algorithms using the SWAX benchmark.

The main contributions of this work are: 1) generation of a public set of real human and wax figure media *in the wild*, indicating that these faces carry most common variations observed in ordinary scenarios and, consequently, represent practical situations; 2) development of straightforward baselines, which stand on long-established techniques in the interest of meeting the requirements of each proposed protocol and guarantee a legitimate comparison among algorithms considering the sort of data they are trained on.

2. RELATED WORKS

Most existing datasets concentrate either on modeling impersonation attack strategies or recapturing genuine images and videos on different mediums. The former demands a high expertise level since it requires the manufacturing of realistic masks and, therefore, is an expensive process. The latter usually involves high-quality cameras for the sake of capturing deceptive print and video replay attacks.



Fig. 2: Face samples collected from online resources and constituting *bona fide* (top row) and *counterfeit* (bottom row) images.

Several face spoofing datasets have been proposed in the last decade [12, 1, 6]. More recently, Boulkenafet *et al.* [7] designed one of the largest mobile-based benchmarks, OULU-NPU, holding more than four thousand real-access and attack videos captured using the frontal camera of six mobile devices. Liu *et al.* [8] introduced SIW, with live and spoof videos of 165 individuals covering a large range of expression, illumination and pose variations. These datasets represent situations that fail to generalize well in conditions where attacks do not proceed from print or replay spoof mediums. More precisely, they require that algorithms learn whether a medium corresponds to an attack rather than the individual itself. A way to make biometric systems robust to other realistic intrusions is to lay out medium-free presentation attacks, like mask-based attacks.

Few masks-based attacks have also been proposed [13, 9]. Manjani *et al.* [11] propose SMAD, a database containing stretchable masks under unconstrained environments, allowing actors to speak and blink, in an attempt to pass on a lifelike sensation. Bhattacharjee *et al.* [3] consider rigid and flexible silicone masks as they release CS-MAD, a dataset captured using low-cost cameras with the collaboration of 14 subjects. The use of masks can restrict natural facial functions/movements such as smiling, talking and blinking, failing to proper fit and match real faces. It becomes even more complex under the need of accurately reproducing another person’s appearance due to the requirement of facial casts. Realistic mask impersonations demand face molds from the subject being modeled and from the person about to wear the disguise. The process of making lifelike masks can cost high figures due to its manufacturing complexity and, to make matters worse, real presentation attacks do not rely on the cooperation of the individual “being cloned”.

There are numerous approaches focusing on print or replay spoofing attacks. Many methods deal with the design of handcrafted descriptors and learning algorithms whereas others focus on the convolutional neural networks trend, described as follows.

There is a myriad of anti-spoofing works in the literature [6, 14, 15, 16]. Valle *et al.* [17] presented a transfer learning method using a pre-trained DNN model on static features to recognize photo, video and mask attacks. Liu *et al.* [8] combined DNN with RNN to estimate the depth of face images along with rPPG signals to boost the detection of unauthorized access. Jourabloo *et al.* [18] decomposed a spoofing face into noise and face information by modeling the process of how the attack is generated from its original live image.

Only recently have researchers turned their attention to fraudulent 3D-mask attacks. Steiner *et al.* [19] adopted spectral signatures of material surfaces in the short wave infrared range to distinguish human skin from other materials, regardless of the hide characteristics. Bhattacharjee *et al.* [20] used off-the-shelf imaging devices to detect print and replay intrusions as well as mask-based presentation attacks near-real-time applications. Many literature methods end up being restricted to specific datasets domains, especially when cameras have comparable capture quality. However, they usually fail to achieve good results on cross-dataset evaluations [14, 21, 8].

3. THE SWAX DATASET

Realistic wax-made sculptures are susceptible to the same limitations found in silicon masks. However, the former tends to be more faithful to actual human traits and, consequently, present greater chances of deceiving biometric systems than masks. They can also be as portable as masks since presentation attacks may require either face portraits or busts instead of full-body figures. From this perspective, we propose the Sense Wax Attack dataset (SWAX)¹.

The SWAX database is compiled from unrestrained online resources and consists of characters, celebrities and public figure images and videos to whom wax dummies have been sculpted into. The database contains 33 female and 22 male individuals. It consists of 1,812 images and 110 videos of 55 people/figures, arranged in training, validation and test sets. More precisely, each subject holds at least 20 authentic still images and a minimum of 10 counterfeit images. All motion and still pictures are manually captured under uncontrolled scenarios, formed by uncooperative individuals and distinct camera viewpoints.

Figure 2 shows the main aspects of the proposed database, depicting authentic and attack samples, respectively. *Bona fide* individual samples also incorporate human appearance alterations due to the natural aging process as pictures are collected without any age restraint. Some subject samples also present accessories such as spectacles or hats, and some male individuals may even have mustaches or beards. Therefore, there is a significant variation in facial expressions, illumination, pose, scene and camera types.

The proposed dataset comprises mutually exclusive sets, which indicates that no subject identity contained in the training set can be made available in the validation or the test sets simultaneously. The same applies to validation and test sets, suggesting that alike subjects should not come out in distinct subsets. For development purposes, the proposed dataset encompasses a leave-one-out cross validation strategy in an attempt to escape unfairly algorithm biases and expose overfitting occurrences. In addition, the data presented in the SWAX dataset should be used as is, according to the provided training, evaluation and test sets.

3.1. Evaluation Protocols

The following paragraphs specify the evaluation protocols guideline and set the proper manipulation of training, validation and test collections. These sets are characterized by the corresponding subject/figure identity in which each face picture² is presented as either *bona fide* (authentic) or *counterfeit* (attack) sample. The four protocols³ are detailed in the following subsections.

¹<http://sense.dcc.ufmg.br/en/dataset/swax-dataset/>

²The word “picture” refers to still images and videos.

³Protocol 01 follows the unsupervised paradigm whereas the remaining ones adhere to the supervised learning task.

3.1.1. Protocol 01: Unsupervised, with additional data

Real-world biometric applications are inclined to anticipate all sorts of illegal intrusions and undergo attacks of distinct nature, which are unknown to the training stage. However, unpredictable attack characteristics require a notable generalization potential that is not usually represented in supervised and multi-class classification techniques as they tend to become too specific to some particular attacks.

It is pertinent to handle spoofing detection problems with one-class-based techniques when labels are not available. Alternate unsupervised strategies comprehend autoencoders [18], clustering [22] and domain adaptation [23] algorithms. One-class classification can be defined as a special case of unsupervised learning where only the class comprising authentic face pictures is well characterized by training data instances. It implies that counterfeit pictures are not known at training time but may emerge at test time. In essence, having all picture samples sharing identical labels in the target variable is analogous to being unlabeled seeing that there is no discriminative information about their class names [24].

For the sake of following the unsupervised protocol appropriately, the following conditions shall apply:

- Procedures claimed to be unsupervised cannot make use of presentation attack samples at the training stage, being restricted to authentic samples only.
- There should be no parameters carrying *bona fide* or *counterfeit* labels, not to mention any other relevant information such as file names or singular identifiers.
- Approaches can neither use “beforehand information” concerning the number of samples included in each class nor use the label distribution in the training and test sets.
- Supplementary samples, outside of SWAX database, are allowed in cases where they depict *bona fide* individuals only and should not be composed of hand-labeled data or information indicating whether pictures are authentic or comprise presentation attacks.

3.1.2. Protocol 02: Restricted, without additional data

The second protocol is in consonance with the supervised learning task, in which there is access to target variables for a set of training data. Each face picture is provided with one out of two categories: *bona fide* or *counterfeit* class. These two labels represent ground-truth information that favor the model in encountering the best-fitted mapping function to guarantee good predictions about “future” picture presentations. That is, authentic and attack samples not listed in the training set to provide more accurate guidance on how well approaches generalize to unseen data.

In contrast to the unsupervised paradigm, this protocol admits information indicating whether a picture is authentic or consists of an attack. Still, it dismisses any type of annotation or data from outside SWAX database, including supplementary picture samples, external tools like facial landmark detectors and alignment methods learned on separate pictures, or feature extractors trained on other data sources. Such external algorithms must be unsupervised, and the data they operate on must be entirely within SWAX database’s training sets. In other words, authentic/attack training labels provided in protocol 02 are authorized to be used along with external algorithms, provided that they satisfy the requirements below:

- Algorithms under this protocol are not allowed to use supplementary data, either to identify presentation attacks or perform any kind of picture preprocessing.
- Validation and test sets are exclusive to their own purposes and cannot be employed to learn auxiliary methods.

- Researchers cannot rely on supplementary labeled data, such as manual face segmentation or facial landmark annotation.

3.1.3. Protocol 03: Unrestricted, with no-label additional data

Differently from first and second protocols, which restrain investigators and developers in such a way they must employ either authentic no-label outside data or SWAX training data alone, respectively; protocol 03 acknowledges the exploitation of additional data sources in the interest of improving an algorithm’s precision. Actually, the main distinction between protocols 02 and 03 is that the latter supports the adoption of SWAX external data if and only if they do not consist of *bona fide* and *counterfeit* labels.

The outside training data may encompass genuine or fraudulent face samples, but the annotation data must be limited to labeled or segmented face components, for instance, mouth contour and eye corners to name a few. Auxiliary algorithms proposed for feature description, face detection and face alignment, even though having the possibility of being tuned on non-SWAX face samples, are legitimate as we understand that outside data and methods used exclusively for no-spoofing reasons is significantly different from using supplementary data/methods to learn spoofing detection classifiers, deserving an exclusive category. Therefore, this protocol grants permission to utilizing feature extractors and alignment algorithms that have been developed for completely unrelated purposes.

The third protocol is distinguished from the others as it is subjected to the following conditions:

- Outside data should not consist of individuals included in SWAX database.
- External pictures cannot hold corresponding information indicating whether they are genuine or fraudulent samples.
- Additional pictures can be labeled with keypoints or segments for the sake of designing preprocessing algorithms.
- Non-SWAX annotations cannot present information that may accredit the formation of authentic or attack face pictures.

3.1.4. Protocol 04: Totally Unrestricted

Supplementary labeled data play an important role in machine learning algorithms as it qualifies biometric algorithms to identify a broader variety of spoofing patterns. On the contrary previously described protocols, protocol 04 endorses the use of SWAX data as well as additional picture samples regardless of the data annotation provided. The totally unrestricted pattern is the most permissive protocol as it admits outside datasets, external feature extractors and other methods that have been built on independent visual data as long as they adhere to the subsequent requirements:

- Supplementary genuine and fraudulent pictures in which their corresponding identity is not available in the SWAX database.
- External face samples may include annotated keypoints, attributes, segments as well as information carrying *bona fide* or *counterfeit* labels.

Protocol 04 admits cross-dataset experiments in the interest of assessing the generalization capability of algorithms and increasing their performance. In consequence, researchers are allowed to avoid SWAX samples, using external data only in the training stage, but compelled to evaluate the designed approaches on the provided testing splits. The idea of combining multiple datasets provides improved statistical power and enhanced classification ability on different domains, turning an algorithm more robust to diversified attacks.

3.2. Evaluation Metrics

Evaluation metrics are employed in order to point out a model’s performance. For the SWAX database, we select two different metrics: the Receiver Operating Characteristic and the standardized ISO/IEC 30107-3 assessment mechanisms for biometric systems.

The IEC 30107-3 designates a particular set of metrics for supervised and unsupervised paradigms [25]. They are denominated Attack Presentation Classification Error Rate, $APCER = \frac{1}{V_{PA}} \sum_{i=1}^{V_{PA}} (1 - Res_i)$; and Bona Fide Presentation Classification Error Rate, $BPCER = \frac{1}{V_{BF}} \sum_{i=1}^{V_{BF}} (Res_i)$. V_{PA} indicates the number of spoofing attacks whereas V_{BF} expresses the total number of authentic presentations. Res_i receives the value 1 when the i -th probe video presentation is categorized as an attack and 0 if labeled as *bona fide* presentation. APCER and BPCER resemble False Acceptance and False Rejection Rates, traditionally employed in the literature when assessing binary classification methods.

The Average Classification Error Rate (ACER) summarizes the overall system performance as it comprehends the mean of APCER and BPCER at the decision threshold determined on the validation set [26]. Although not used in this work, researchers can also consider the widely explored Detection Error Trade-off curve (DET) [27], a graphical plot of error rates (comprising missed detections and false alarms) usually engaged in the evaluation of binary classification systems for unsupervised tasks.

4. BASELINE EXPERIMENTS

This section analyzes the performance of some methods on the SWAX benchmark and also details the employed feature descriptors, additional dataset and straightforward spoofing baselines.

External Datasets. Protocol 01 and 04 authorize researchers to incorporate external data. We select LFW [28] database, initially designed for studying the problem of unconstrained face recognition, to compose the unsupervised analysis. LFW is compatible with SWAX’s first protocol as it only comprises face images of distinct public figures, having no incidence of waxwork models. MFSD [6] satisfies the totally unrestricted protocol requirements as it provides authentic/attack label information for the training stage.

Feature Descriptors. We select GLCM [29] and LBP [30] for the extraction of visual features in lower dimensional spaces. GLCM features are computed with directions $\theta \in \{0, 45, 90, 135\}$ degrees, distance $d \in \{1, 2\}$, 16 bins and six texture properties: contrast, dissimilarity, homogeneity, energy, correlation, and angular second moment. They are computed from each image’s corresponding Fourier transform in the frequency domain [14]. LBP information comprises 256 bins, a radius equal to 1, and eight points arranged in a 3×3 matrix. They are taken from each image’s color band after being converted from RGB into HSV and $Y_C R_C B_C$ [31]. ALL refers to experiments in which GLCM and LBP descriptors are extracted and concatenated.

Baselines. As showed in Table 1, some approaches have been picked out from literature and employed as baselines, meeting the requirements and cross validation fold splits of each SWAX’s protocol. WSVM [32] is a classification algorithm that provides solutions for non-linear classification in open-set scenarios. For the unsupervised paradigm, we combine one-class WSVM with different feature descriptors in an attempt to encounter an adequate boundary between authentic and attack samples.

Table 1: APCER and BPCER results (%) on SWAX’s protocols. Supplementary datasets are depicted between parenthesis.

Prot.	Method	APCER	BPCER	ACER
1	GLCM+WSVM	11.9 ± 9.46	86.3 ± 9.27	49.1 ± 9.37
	LBP+WSVM	82.3 ± 6.65	10.1 ± 4.68	46.2 ± 5.66
	ALL+WSVM	90.9 ± 5.07	10.9 ± 4.93	50.9 ± 5.00
	ALL+WSVM(LFW)	90.6 ± 5.09	09.8 ± 5.03	50.2 ± 5.06
2	RHYTHM	37.6 ± 6.66	12.3 ± 4.59	24.9 ± 5.63
	EPLS	35.7 ± 9.35	4.34 ± 2.23	20.1 ± 6.29
	ESVM	48.3 ± 9.63	4.87 ± 2.03	26.5 ± 6.33
3	ALIGN+RHYTHM	32.4 ± 7.19	13.4 ± 4.79	22.9 ± 5.99
	ALIGN+EPLS	33.3 ± 9.75	4.67 ± 2.28	18.9 ± 6.01
	ALIGN+ESVM	45.9 ± 9.24	4.64 ± 2.01	25.3 ± 6.13
4	DE-SPOOFING	12.1 ± 5.52	60.2 ± 10.5	36.1 ± 8.05
	EPLS(MFSD)	28.2 ± 6.40	5.56 ± 1.97	16.8 ± 4.19
	ESVM(MFSD)	39.8 ± 5.55	5.44 ± 1.52	22.6 ± 4.53

RHYTHM [14] searches for frequency domain artifacts as it is trained and tested on SWAX’s second and third protocols. Since Protocol 03 admits external data and methods as long as they are not designed for spoofing purposes, additional algorithms can be added in preprocessing tasks. ALIGN [33] is a DNN-based face alignment method that iteratively improves the locations of the facial landmarks by positioning subject faces into a canonical pose. For Protocol 04, DE-SPOOFING [18] performs a cross-dataset evaluation as it is trained on OULU-NPU dataset [7] and then estimate deep spoof noises from SWAX probe face samples. EPLS and ESVM methods [16] consist of an ensemble of classifiers learned on random subsets of the training set. At test time the method projects a probe image onto all trained classifiers and computes the ratio of the number of positive responses attained to the total number of classification models. If most classifiers return positive responses, the face image is likely to be an authentic sample, or a spoofing attack, otherwise.

Evaluated algorithms provide an insight into the performance of countermeasure methods when confronted with wax figure images. Experiments in Table 1 show that advanced spoofing detection methods are still vulnerable to high-quality violation attempts, demanding researchers to deliver more accurate biometric systems. The attained results are not satisfactory and reveal that the face anti-spoofing problem has a large room for improvement despite the significant progress in recent years.

5. CONCLUSIONS

This work comprises a public benchmark of real human and wax figure images and videos, entitled Sense Wax Attack dataset, developed for face anti-spoofing purposes. The database consists of different protocols, cross validation splits and evaluation metrics to operate fair comparisons among different algorithms. Results show that wax figures can be employed in pursuance of bypassing biometric systems and obtain illegitimate access. Therefore, there is a lot to improve when it comes to wax-based face spoofing attacks. We hope that with the advent of this dataset researchers will consider new sorts of spoofing attacks as a motivation to delivering robust methods to the challenging area of spoofing detection.

6. ACKNOWLEDGMENTS

The authors would like to thank the National Council for Scientific and Technological Development – CNPq (Grants 438629/2018-3 and 309953/2019-7), the Minas Gerais Research Foundation – FAPEMIG (Grants APQ-00567-14 and PPM-00540-17), and the Coordination for the Improvement of Higher Education Personnel – CAPES (DeepEyes Project). This study was financed in part by the Coordenacao de Aperfeiçoamento de Pessoal de Nivel Superior - Brasil (CAPES) - Finance Code 001.

7. REFERENCES

- [1] Ivana Chingovska, André Anjos, and Sébastien Marcel, "On the effectiveness of local binary patterns in face anti-spoofing," in *International Conference of Biometrics Special Interest Group*, 2012.
- [2] Allan Pinto, Helio Pedrini, William Robson Schwartz, and Anderson Rocha, "Face spoofing detection through visual codebooks of spectral temporal cubes," *Transactions on Image Processing*, vol. 24, no. 12, 2015.
- [3] Sushil Bhattacharjee, Amir Mohammadi, and Sébastien Marcel, "Spoofing deep face recognition with custom silicone masks," in *Conference on Biometrics Theory, Applications and Systems*. IEEE, 2018.
- [4] Sandeep Kumar, Sukhwinder Singh, and Jagdish Kumar, "A comparative study on face spoofing attacks," in *International Conference on Computing, Communication and Automation*. IEEE, 2017.
- [5] Shenzhen Daily, "6 arrested for pyramid sales scam," 2012.
- [6] Di Wen, Hu Han, and Anil K Jain, "Face spoof detection with image distortion analysis," *Transactions on Information Forensics and Security*, vol. 10, no. 4, 2015.
- [7] Zinelabidine Boulkenafet, Jukka Komulainen, Lei Li, Xiaoyi Feng, and Abdenour Hadid, "Oulu-npu: A mobile face presentation attack database with real-world variations," in *International Conference on Automatic Face & Gesture Recognition*. IEEE, 2017.
- [8] Yaojie Liu, Amin Jourabloo, and Xiaoming Liu, "Learning deep models for face anti-spoofing: Binary or auxiliary supervision," in *Conference on Computer Vision and Pattern Recognition*. IEEE, 2018.
- [9] Siqi Liu, Baoyao Yang, Pong C Yuen, and Guoying Zhao, "A 3d mask face anti-spoofing database with real world variations," in *Conference on Computer Vision and Pattern Recognition Workshops*. IEEE, 2016.
- [10] Akshay Agarwal, Daksha Yadav, Naman Kohli, Richa Singh, Mayank Vatsa, and Afzel Noore, "Face presentation attack with latex masks in multispectral videos," in *Conference on Computer Vision and Pattern Recognition Workshops*. IEEE, 2017.
- [11] Ishan Manjani, Snigdha Tariyal, Mayank Vatsa, Richa Singh, and Angshul Majumdar, "Detecting silicone mask-based presentation attack via deep dictionary learning," *Transactions on Information Forensics and Security*, vol. 12, no. 7, 2017.
- [12] Zhiwei Zhang, Junjie Yan, Sifei Liu, Zhen Lei, Dong Yi, and Stan Z Li, "A face antispoofing database with diverse attacks," in *International Conference on Biometrics*. IEEE, 2012.
- [13] Nesli Erdogmus and Sébastien Marcel, "Spoofing in 2d face recognition with 3d masks and anti-spoofing with kinect," in *Conference on Biometrics Theory, Applications and Systems*. IEEE, 2013.
- [14] Allan Pinto, William Robson Schwartz, Helio Pedrini, and Anderson de Rezende Rocha, "Using visual rhythms for detecting video-based facial spoof attacks," *Transactions on Information Forensics and Security*, vol. 10, no. 5, 2015.
- [15] Lei Li, Xiaoyi Feng, Zinelabidine Boulkenafet, Zhaoqiang Xia, Mingming Li, and Abdenour Hadid, "An original face anti-spoofing approach using partial convolutional neural network," in *International Conference on Image Processing Theory, Tools and Applications*. IEEE, 2016.
- [16] Rafael Henrique Vareto, Matheus A Diniz, and William Robson Schwartz, "Face spoofing detection on low-power devices using embeddings with spatial and frequency-based descriptors," in *Iberoamerican Congress on Pattern Recognition*. Springer, 2019, pp. 187–197.
- [17] Eduardo Valle and Roberto Lotufo, "Transfer learning using convolutional neural networks for face anti-spoofing," in *International Conference Image Analysis and Recognition*. Springer, 2017.
- [18] Amin Jourabloo, Yaojie Liu, and Xiaoming Liu, "Face de-spoofing: Anti-spoofing via noise modeling," in *European Conference on Computer Vision*. Springer, 2018.
- [19] Holger Steiner, Sebastian Sporrer, Andreas Kolb, and Norbert Jung, "Design of an active multispectral swir camera system for skin detection and face verification," *Journal of Sensors*, vol. 2016, 2016.
- [20] Sushil Bhattacharjee and Sébastien Marcel, "What you can't see can help you-extended-range imaging for 3d-mask pad," in *International Conference of Biometrics Special Interest Group*. IEEE, 2017.
- [21] Zinelabidine Boulkenafet, Jukka Komulainen, and Abdenour Hadid, "Face spoofing detection using colour texture analysis," *Transactions on Information Forensics and Security*, vol. 11, no. 8, 2016.
- [22] Olegs Nikisins, Amir Mohammadi, André Anjos, and Sébastien Marcel, "On effectiveness of anomaly detection approaches against unseen presentation attacks in face anti-spoofing," in *International Conference on Biometrics*. IEEE, 2018.
- [23] Haoliang Li, Wen Li, Hong Cao, Shiqi Wang, Feiyue Huang, and Alex C Kot, "Unsupervised domain adaptation for face anti-spoofing," *Transactions on Information Forensics and Security*, vol. 13, no. 7, 2018.
- [24] Wei Liu, Gang Hua, and John R Smith, "Unsupervised one-class learning for automatic outlier removal," in *Conference on Computer Vision and Pattern Recognition*. IEEE, 2014.
- [25] "Information technology - biometric presentation attack detection," Tech. Rep., ISO/IEC JTC 1/SC 37 Biometrics, 2017.
- [26] Z Boulkenafet, J Komulainen, Z Akhtar, A Benlamoudi, D Samai, SE Bekhouche, A Ouafi, F Dornaika, A Taleb-Ahmed, L Qin, et al., "A competition on generalized software-based face presentation attack detection in mobile scenarios," in *International Joint Conference on Biometrics*. IEEE, 2017.
- [27] Alvin Martin, George Doddington, Terri Kamm, Mark Ordowski, and Mark Przybocki, "The det curve in assessment of detection task performance," Tech. Rep., National Inst of Standards and Technology Gaithersburg MD, 1997.
- [28] Gary B Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller, "Labeled faces in the wild: A database for studying face recognition in unconstrained environments," in *Workshop on Faces in 'Real-Life' Images: Detection, Alignment, and Recognition*, 2008.
- [29] Robert M Haralick, Karthikeyan Shanmugam, and Its' Hak Dinstein, "Textural features for image classification," *Transactions on systems, man, and cybernetics*, , no. 6, pp. 610–621, 1973.
- [30] Audrey Ledoux, Olivier Losson, and Ludovic Macaire, "Color local binary patterns: compact descriptors for texture classification," *Journal of Electronic Imaging*, vol. 25, no. 6, pp. 061404, 2016.
- [31] Zinelabidine Boulkenafet, Jukka Komulainen, and Abdenour Hadid, "Face anti-spoofing based on color texture analysis," in *International Conference on Image Processing*. IEEE, 2015.
- [32] Walter J Scheirer, Lalit P Jain, and Terrance E Boult, "Probability models for open set recognition," *Transactions on Pattern Analysis and Machine Intelligence*, vol. 36, no. 11, 2014.
- [33] Marek Kowalski, Jacek Naruniec, and Tomasz Trzcinski, "Deep alignment network: A convolutional neural network for robust face alignment," in *Conference on Computer Vision and Pattern Recognition Workshops*. IEEE, 2017.