

 Open access • Book Chapter • DOI:10.1007/978-3-319-28501-6_8

Face Recognition Systems Under Spoofing Attacks — Source link

Ivana Chingovska, Nesli Erdogmus, André Anjos, Sébastien Marcel

Institutions: Idiap Research Institute

Published on: 01 Jan 2016

Topics: Spoofing attack

Related papers:

- [Face Spoof Detection With Image Distortion Analysis](#)
- [A face antispoofing database with diverse attacks](#)
- [Spoofing in 2D face recognition with 3D masks and anti-spoofing with Kinect](#)
- [On the effectiveness of local binary patterns in face anti-spoofing](#)
- [OULU-NPU: A Mobile Face Presentation Attack Database with Real-World Variations](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/face-recognition-systems-under-spoofing-attacks-28j93gca33>



**FACE RECOGNITION SYSTEMS UNDER
SPOOFING ATTACKS**

Ivana Chingovska

Nesli Erdoganmus

André Anjos

Sébastien Marcel

Idiap-RR-18-2020

SEPTEMBER 2020

Face recognition Systems Under Spoofing Attacks

Ivana Chingovska, Nesli Erdogmus, André Anjos and Sébastien Marcel

Abstract In this chapter we give an overview of spoofing attacks and spoofing counter-measures for face recognition systems, with a focus on Visual Spectrum systems (VIS) in 2D and 3D, as well as Near Infrared (NIR) and multispectral systems. We cover the existing types of spoofing attacks and report on their success to bypass several state-of-the-art face recognition systems. The results on two different face spoofing databases in VIS and one newly developed face spoofing database in NIR, show that spoofing attacks present a significant security risk for face recognition systems in any part of the spectrum. The risk is partially reduced when using multispectral systems. We also give a systematic overview of the existing anti-spoofing techniques, with an analysis of their advantages and limitations and prospectives for future work.

1 Introduction

Thanks to the growing availability of inexpensive cameras, as well as the unobtrusiveness of capturing procedures, face has a guaranteed position as one of the most exploitable biometric modes. Its wide deployment is further reinforced by the rapid advancement of face recognition systems, which nowadays provide reliable recognition even under challenging conditions. Historically, 2D face recognition in the

Ivana Chingovska

Idiap Research Institute, Switzerland e-mail: ivana.chingovska@idiap.ch

Nesli Erdogmus

IZTECH Dept. of Computer Engineering, Turkey e-mail: neslierdogmus@iyte.edu.tr

André Anjos

Idiap Research Institute, Switzerland e-mail: andre.anjos@idiap.ch

Sébastien Marcel

Idiap Research Institute, Switzerland e-mail: sebastien.marcel@idiap.ch

visible spectrum (VIS) has got the most attention and has reached a stage where it provides a secure, robust and trustworthy biometric authentication at different security checkpoints: ID control systems, protected web services, and even mobile devices. On the other hand, face recognition in 3D, near-infrared (NIR) and thermal spectrum shows an increased popularity in the recent years [21, 65].

Unfortunately, face recognition systems can be an attractive target for spoofing attacks: attempts to illegally access the system by providing a copy of a legal user's face. Information globalization acts in favor of such system misuse: users' personal data, including face images and videos, are nowadays widely available and can be easily downloaded from the Internet. Printed photographs of a user face, digital photographs displayed on a device, video replays and 3D masks have already proven to be a serious threat for face recognition systems in VIS. Spoofing attacks for NIR face recognition systems have not received as much attention, but recent spoofing attempts indicate on their vulnerability too [78]. Thermal imaging face recognition systems, which use the thermal emissions pattern naturally coming from the human body [65] are naturally resistant to any kind of 2D spoofing attacks, as the thermal signatures of 2 individuals are different [56]. Even more, such systems are resistant to surgically performed face alterations: because tissue redistribution, addition of artificial materials, as well as alteration of blood vessel flows that may happen during a surgery have a big impact on the thermal signature of a person [54].

In this chapter, we will cover research attempts in spoofing and anti-spoofing for the face mode from two perspectives. Firstly, we will investigate to what extent the state-of-the-art face recognition systems are vulnerable to spoofing attacks. This is a vital step towards verifying the threat and justifying the need of anti-spoofing methods. In addition, this step may reveal whether a spoofing attack database is relevant to be used to develop and evaluate anti-spoofing methods. We perform this analysis on four state-of-the-art face recognition systems working in VIS and NIR. In VIS, we exploit two different publicly available face spoofing databases, one with 2D attacks, and one with 3D mask attacks. To perform the analysis in NIR, we develop and present the first publicly available face spoofing database containing VIS and NIR spoofing attacks. By fusing the scores of the systems working in VIS and NIR, we extend the analysis to multispectral systems as well.

Secondly, we give an overview of the recent advancements in counter-measures to spoofing attacks for face recognition systems. This includes systematic categorization of the anti-spoofing methods and investigation on the attacks they are effective against. While there is a plethora of anti-spoofing methods for VIS face recognition systems, the amount of methods for NIR and multispectral systems is significantly smaller.

Unfortunately, it is extremely difficult to comparatively evaluate the performance of the existing anti-spoofing methods, mainly due to two factors. Firstly, very few of the research papers release the source code and the exact parameters to reproduce the presented results. Secondly, many of them are evaluated on private databases, or are targeting just one type of spoofing attacks. Therefore, while we most often omit performance numbers, we distinguish methods whose results are fully reproducible on publicly available databases.

This chapter is organized as follows. We cover 2D face recognition systems in VIS and NIR under 2D spoofing attacks in Sections 2.1 and 2.2, respectively. In Section 3, we cover face recognition systems in VIS under 3D spoofing attacks. Conclusions and discussion follow in Section 4.

2 Face recognition systems under 2D spoofing attacks

2.1 *Visual spectrum (VIS) face recognition systems*

Numerous spoofing attack trials to test the robustness of commercial devices [17, 42], as well as several face spoofing databases have proved that face recognition systems in VIS can be spoofed with many different types of attacks. The attacks differ by their complexity, their cost and the amount of effort and skills required for producing them. The effectiveness of the attacks is closely related with these properties.

The spoofing counter-measures developed to protect 2D face recognition systems in VIS are by now developed to a very good extent. For example, the 2nd competition of counter-measures to 2D face spoofing attacks [14], where two of the submitted algorithms achieved perfect spoofing detection rate. The objective of this Section is to summarize the research efforts in this direction, in terms of available spoofing attack types and databases, as well as existing solutions. We focus on face verification systems, where the spoofing attacks make most sense.

2.1.1 Types of attacks and databases

Probably the most simple type of face spoofing attack is the print attack, which consists of printing a photo of a valid user's face on paper. A more sophisticated type of attack involves presenting a digital photo on the screen of a mobile device. These two types of attacks retain the face appearance, but present only a static face showing no signs of vitality. More sophisticated versions of the printed attacks simulate vitality by perforating the eye region or moving, rotating and warping the printed paper [66, 52, 83]. In addition, there are video replay spoofing attacks, where a face video of a valid user is presented on the screen of a mobile device. Examples of spoofing attacks based on drawing of a user's face or using make-up to masquerade as a valid user have been registered at the ICB 2013 spoofing challenge¹. Attacks with 3D masks will be covered in Section 3.

Besides the way of reproducing the spoofed face, the spoofing attacks can differ in a number of other criteria. For example, they can be recorded in controlled or uncontrolled environments. Furthermore, a fixed or a hand support can be used

¹ <http://www.biometrics-center.ch/testing/tabula-rasa-spoofing-challenge-2013>

for holding the spoof medium [12]. [34] defines the term scenic and close-up 2D spoofing attack. Scenic attacks refers to attacks where the background content of the original image used to create the spoofing attack is visible alongside the spoofed face. On close-up attacks, the border of the spoofing medium is fully visible. Although the close-up attacks contain an additional cue for detection, they are inevitable in cases where the size of the spoofing mediaum is too small.

The available face spoofing databases cover a different subsets of these types of attacks. Different types of spoofing attacks pose a different level of difficulty to detect and are usually addressed with different types of counter-measures. The number of face spoofing databases which are publicly available is limited. Up to the present moment, the established counter-measures to 2D face spoofing attacks have been evaluated either on private databases, or on three publicly available face spoofing databases: NUAA Photograph Imposter Database [66], CASIA Face Anti-spoofing Database (CASIA-FASD) [83] and the Replay-Attack family of databases [12]. NUAA database consists of attacks with printed photographs. It contains still images of real accesses and attacks to 15 identities and is recorded in three sessions under three different illumination conditions. When capturing the attacks, the photographs of the users are translated, rotated or warped. The attacks in NUAA database, few of which are shown in Fig. 1, are examples of close-up spoofing attacks.



Fig. 1 Real access (first column) and examples of attack samples from NUAA database

CASIA-FASD provides videos of real accesses and three types of attacks to 50 identities. The first type is performed with printed photographs warped in front of the camera. The second type are printed photographs with perforated eye regions, so that a person can blink behind the photograph. The third type is a video playback of the user. The two types of printed attacks are taken in a close-up fashion, while the video attacks are scenic attacks. When recording the database, three imaging qualities are considered: low, normal and high, depending on the device used to capture the attacks. Examples of the three types of spoofing attacks in CASIA-FASD are shown in Figure 2.

The Replay-Attack family of databases consists of Print-Attack [5] containing printed photographs, Photo-Attack [4] containing printed and digital photographs, and Replay-Attack [12], as a superset of the previous two databases to which video attacks have been added. There is a total of 50 identities, recorded in both controlled and uncontrolled conditions, with diverse acquisition equipment.

Not all of the spoofing databases have equally wide applicability for evaluating anti-spoofing systems. For example, a database which offers still images, like



Fig. 2 Real access (first column) and examples of attack samples from CASIA-FASD database



Fig. 3 Real access (first column) and examples of attack samples from Replay-Attack database. First row show samples taken in controlled, and second row samples taken in adverse conditions.

NUAA, can not be used for evaluation of counter-measures which require video inputs, like the motion-based algorithm described in Section 2.1.3. In addition, some databases are lacking a protocol to precisely define training, development and test set. Finally, as described in Section 2.1.2, spoofing databases should provide enrollment samples which can be used to train and evaluate a baseline face verification system [13]. Both NUAA and CASIA-FASD suffer from this last drawback, and hence their effectiveness in bypassing face verification systems can not be properly evaluated. This disadvantage is overcome by the databases of the Replay-Attack family.

2.1.2 Assessing the vulnerability


When evaluating a face verification system, it is a common practice to report False Acceptance Rate (FAR) (or False Match Rate (FMR)) and False Rejection Rate

(FRR) (or False Non-Match Rate (FNMR))². The error rate at the point where these two values are equal is called Equal Error Rate (EER), while their average is called Half Total Error Rate (HTER). If the systems are exposed to spoofing attacks, their vulnerability is usually measured using Spoof False Acceptance Rate (SFAR) [13]. If the face verification system is tuned to work at particular operating point (decision threshold), SFAR gives the ratio of spoofing attacks whose score is higher than that point and are thus accepted by the system.

In order to be used for evaluation of verification systems, spoofing attack databases need to have properties that allow for their training [13]. In particular, they need to contain enrollment samples used to enroll clients in the verification systems. Out of the publicly available 2D face spoofing databases, only the Replay-Attack family satisfies this property. Using Replay-Attack database, we trained face verification system based on Gaussian Mixture Model (GMM), which extracts Discrete Cosine Transform (DCT) features from the input images [9]. Fig. 4 shows the distribution of the scores for the real accesses, zero-effort impostors and spoofing attacks from Replay-Attack for this system. The green line depicts the point which is chosen as a decision threshold based on EER criteria depending on FAR and FRR. The system shows a remarkable separability between the score distributions of the real accesses and zero-effort impostors, resulting in an almost perfect verification results (HTER=0.14%). However, the distributions of the scores of the real accesses and spoofing attacks overlap by a large extent. As a result, the system accepts 91.5% of the spoofing attacks, which proves its high vulnerability to spoofing.

We performed similar analysis for three additional state-of-the-art face verification systems, each of which is based on different features and modeling paradigm. The first one uses Local Gabor Binary Pattern Histogram Sequences (LGBPHS) [79], the second one is based on Gabor Jets comparison (GJet) [27], while the third one uses Inter-Session Variability modeling (ISV) [70]. The results are shown in Table 1. All of the examined systems perform very well in the verification task. However, with SFAR of 90%, each one of them exhibits a high vulnerability to spoofing, demonstrating the need for development of suitable counter-measures.

Fig. 4 Score distribution of GMM based face recognition system for the samples in Replay-Attack. Real accesses: ■, zero-effort impostors: ■, spoofing attacks: ■



² In their formal definition, FAR and FMR and FRR and FNRM are not synonymous [47]. However, they can be treated as such in some special cases, and we will do so, following the practice adopted in [28].

Table 1 Verification error rates and spoofing vulnerability of baseline face verification systems (in %)

system	FAR	FRR	SFAR
GMM	0.05	0.24	91.5
LGBPHS	1.47	2.13	88.5
GJet	0.28	0.24	95.0
ISV	0.00	0.17	92.6

2.1.3 Spoofing counter-measures

The anti-spoofing methods for the face mode can be primarily categorized based on the type of data that is used to detect the spoofing attacks. In this respect, they can fall into two categories: hardware-based and software-based [63]. The hardware-based solutions use additional hardware to detect the spoofing attacks, which may be a thermal or NIR camera, 3D sensor etc. The software-based ones utilize solely the information which is captured by the camera of the recognition system and try to directly exploit the characteristic of the input images.

Some of the software-based methods require, either implicitly or explicitly, that the user answers to some kind of interactive challenge. Yet, most of these methods take the decision in a non-intrusive manner, without any requirement for an explicit input from the user. They use different types of cues that may indicate the presence of a live subject in front of the system: liveness, motion, visual appearance, contextual information and 3D reconstruction information. Usually, the features extracted for these purposes are hand-crafted based on prior knowledge about the task; however, there are algorithms which extract relevant features in a completely data-driven fashion.

In the remainder of this Section, we are going to cover the most prominent representatives of face anti-spoofing methods, and make a comparative analysis of their performance and limitations. We will put an additional note to those which depend on interaction with the user.

Before proceeding, it is important to notice that several researchers have made attempts to increase the robustness of biometric recognition systems to spoofing attacks by using multiple biometric modes [62]. The intuition behind these solutions is that an attacker may need more effort to spoof the system, because she needs to spoof more modes. Within such multimodal framework, face has been combined with fingerprint and iris [29, 1, 61, 60], or with voice [10]. [29, 1, 61, 60] have proven, however, that poorly designed combination rules for multimodal systems may not be helpful. Combination rules designed specifically for the purpose of increased robustness have been designed in [61, 60].

Liveness detection

The liveness detection anti-spoofing methods base their decision on the evidence of liveness present on the scene. Usually, eye-blinking, mouth movements and in-

voluntary subtle head movements are considered as evidence of liveness. One of the first attempts to employ eye-blinking for anti-spoofing is performed by [51], which uses Conditional Random Fields (CRF) to model the state of the eye as open or closed and the correlation between its state and the observation. With a similar purpose, [71] uses active shape models to detect the eye contours and difference of images to detect the blinking activity. In [32], eye-blinking detection is combined with analysis of the 3D properties of the subject.

A key, but limiting assumption of the liveness detection methods is that the subject will experience the actions that suggest liveness within a given short time frame. For example, [51] assumes that eye blinks happen every 2-4 seconds, which may not be true always and for all the subjects. To be fully successful, these methods depend on user input like deliberate eye-blinks, which may give them a level of intrusiveness.

An attempt to overcome this limitation is done by methods which rely on more subtle changes in the face region, including color changes due to blood flow. To be able to detect these changes, [8] performs Eulerian motion magnification [74] as a pre-processing before applying a technique for analyzing the texture or the motion patterns.

Another drawback of the liveness methods is that, although they may be successful in the case of print and attacks (even when they are warped or rotated [51]), they may be easily deceived by spoofing attacks where liveness evidence is present, like video playback or 3D masks.

Motion analysis

The motion based methods try to find properties of the motion patterns of a person in front of the system, in order to distinguish them from motion patterns in presence of a spoofing attack. A few of these methods base their approach on the assumption that a person's head, being a 3D object, moves differently than a 2D spoofing attack displayed on a planar media. For example, [33] uses optical flow method to track movements on different face parts. The authors assume that, in contrast to a face displayed on a 2D surface, a 3D face will generate higher amount of motion in central face parts closer to the camera (like the nose) than in the face parts which are further away from the camera (like the ears). Furthermore, a 3D face exhibits motion flows which are in opposite directions for central and peripheral face parts. On the other hand, [7] derives a heuristics for the optical flow field for four basic 2D surface motion types: translation, in-plane rotation, panning and swing. On the contrary, a 3D face and facial expressions generate irregular optical flow field.

Another set of motion-based methods assumes a high correlation between the movements in the face region and the background in the case of a spoofing attack. Such a correlation is unlikely in the case of a real access. [5] bases the computation of the correlation on 10 quantities extracted from the face region and the background. For the same purpose, [4] relies on quantization of optical flow motion vectors, while [75] performs foreground-background consistency analysis.

Similarly to the liveness detection methods, the motion analysis approaches depend on the subtle involuntary movements of the user. In addition, sometimes they capture the motion introduced by an attacker who holds the attack media with his hands. If the presumed motion patterns are absent during the short acquisition process (for example, a very still person who does not blink), the methods may fail. These methods are mostly targeting photo spoofing attacks and will most likely fail in case of spoofing attacks by video playbacks or 3D masks. Furthermore, the methods based on motion correlation are particularly directed for scenic 2D spoofing attack, where the background of the spoofed image is visible.

Visual appearance

The anti-spoofing methods analyzing the visual appearance stand behind a strong argumentation about the differences in the visual properties of real accesses and spoofing attacks, explained in a number of publications. Firstly, a real face and the human skin have their own optical qualities (absorption, reflection, scattering, refraction), which other materials that can be used as spoofing media (paper, photographic paper or electronic display) do not possess [53]. Similar differences can appear as a result of the diffuse reflection due to a non-natural shape of the spoofing attacks [77]. Limited resolution of the device used for spoofing or the involuntary shaking of the spoofing media may cause a blurring in the case of spoofing attacks [43, 46, 77]. Artifacts appearing in the spoofing production process, like jitter and banding in the case of print attacks [46, 75] or flickering and Moiré effect in the case of video attacks [55] are yet another sources of differences between the real accesses and spoofing attacks. Many of these visual properties are indistinguishable for the human eye, but often can be easily extracted using different image processing and computer vision algorithms.

The first approach leveraging on the argument that spoofing attacks are usually of lower resolution and thus contain less high-frequency components is proposed in [43]. The proposed feature vector is based on analysis of the 2D Fourier spectrum of the input image and its energy change over time. Instead of comparing the high-frequency content of the input, [66] and [83] base their discrimination on the high-middle band of the Fourier spectrum, which is extracted using Difference of Gaussians (DoG) method.

Some publications assume that the differences between real accesses and attacks are most prominent within the reflectance component of the input image and estimate it in different ways: [66] uses the Lambertian reflectance model [50] and Variational Retinex-based method, while [6] uses dichromatic reflection model. Then, [66] classifies the obtained features using Sparse Low-rank bilinear discriminative model, while [6] compares the gradient histograms of the reflectance images.

A feature set inspired by a physics-based model for recaptured images, which reveals differences in the background contextual information, reflection, surface gradient, color, contrast, chromaticity and blurriness, is created by [26]. Different sets of visual features related to texture, color, edges and/or gradient are used

by [67, 64]. [25] generalizes the appearance differences into quality differences and uses a feature vector composed of 25 different image quality measures.

Several publications make use of specific computer vision descriptors for texture analysis. Local Binary Pattern (LBP) [49] appears to be the most significantly exploited for the purpose of anti-spoofing, both in its single resolution [12] and multiresolution [45, 46, 77] variants. Histogram of Oriented Gradients (HOG) [64, 46, 77], Grey-level Co-occurrence Matrix (GLCM) [64], Haar wavelets [75] and Gabor wavelets [46] are some of the other alternatives.

More recently, the analysis of the visual appearance has been enhanced into a temporal domain. In [55], the authors firstly extract the noise from each video frame, and then summarize the relevant components of its 2D Fourier analysis into so-called Visual Rhythm image. The properties of this image are then captured using GLCM. The method proposed in [23] utilizes LBP-TOP [82], where instead of LBP analysis on a single frame, dynamical LBP analysis on a frame and its neighboring frames is performed.

The methods described before present different rates of success, which can not be easily compared because they are obtained on different types of attacks and usually on databases which are not released publicly. An interesting property of the majority of the visual appearance methods is that they can work even if only a single image is available at input. They are usually applied either on the face bounding box, face parts, or on the full input image. As one of their advantages, they are very user-friendly and non-intrusive and do not depend on the behavior of the user (unlike the liveness detection and motion analysis methods). Furthermore, an attack which can deceive them a-priori has not been presented up to this moment. For example, they can be expected to successfully detect print, photo, video or even 3D mask attacks. Yet, their success may be put into question if the spoofing attacks are printed or displayed on high-resolution media, thus lacking some of the artifacts that these methods rely on. Their generalization properties when applied to different acquisition conditions or new types of attacks are also uncertain, since the visual appearance of the input images often depends on the light condition, acquisition devices or display media.

Contextual information

The context of the scene present as a background information in front of the recognition system is used as a cue to detect spoofing attacks. In [52], the authors notice that in the case of a spoofing attack, there will be a change in the contextual information of the background when the face appears. To detect such changes, the authors compare the regions around reference fiducial key points in the region around the face.

The approach presented in [35] is targeting attacks where the contextual information consists of the border of the spoofing medium. Hence, a prerequisite is that the spoofing medium is fully visible to the system. The method relies on Histogram of Oriented Gradients (HOG) [15] to detect upper body and spoofing medium borders.

3D information

The 3D property of a human face is a cue that unambiguously distinguishes real accesses from 2D spoofing attacks. This is used by several publications, which try to reconstruct or estimate the 3D information from the user's face. For example, [72] recovers and classifies the 3D structure of the face based on two or more images taken from different viewing angles. For similar purposes, [16] uses 3D projective invariants of a moving head. The disadvantage of these approaches is their intrusiveness: the user needs to be collaborative and move his head to a different angle in the first case, or perform certain movements at random intervals in the second case. Avoiding such a constraint, [30] estimates the focus variabilities between two images taken consecutively and focused on different parts of the face. In the case of a 2D spoofing attacks, it is expected that focus variabilities will be absent.

It is important to note that the success of this set of methods is usually limited to 2D spoofing attacks, and are likely to fail 3D mask attacks.

Challenge-response

Unlike the majority of motion analysis of liveness detection methods which rely on the involuntary movements of the user, challenge-response anti-spoofing methods explicitly ask the user to perform certain action to verify his liveness. Representatives of this type have been already mentioned [72, 16]. There are various types of challenges that a user can perform: taking a particular head pose [24] or following a moving point with a gaze [2] are some of them. Finding the static and dynamic relationship between face and voice information from a speaking face or modeling a speaker in 3D shape is an option for anti-spoofing in a multimodal audio-visual system [11]. It is important to note that the last approach can successfully detect not only visual, but even audio-visual spoofing attacks, like video playbacks with recorded utterance or 3D synthetic talking heads.

The challenge-response methods are considered to be intrusive, non-friendly and uncomfortable from the aspect of a user experience. In addition, they usually require that the authentication is performed during a prolonged time span. Finally, they are not transparent for the user. In this way, it is possible for a malicious user to guess the liveness cue and try to bypass it.

Feature learning

Following a recent trend, the anti-spoofing community started experimenting with approaches where the anti-spoofing features are automatically learned directly from the data. This is in contrast to the previously discussed approaches, where the features are inspired by some particular characteristics that can be observed as common either for real accesses or for some type of spoofing attacks. It is argued, however, that the features engineered in this way are not suitable for different kinds of spoof-

ing attacks [76, 48]. Both [76] and [48] are training a Convolutional Neural Network (CNN) for the purpose. In [76] experiments with face images in 5 different resolutions are given, while in [48], the authors use an optimization procedure to select the best CNN to learn the features, out of a family of CNNs with different hyper-parameters.

Hardware-based methods

The hardware-based methods employ an additional piece of hardware along the camera used by the recognition system. These methods detect spoofing attacks using the cues captured by the additional hardware. For example, very often, these methods exploit the properties of the human body in different regions of the electromagnetic spectrum. In such a case, the additional hardware may refer to the sensor used to capture data at a particular wavelength, a light filter which is applied to the camera, or illuminator emitting light at a particular wavelength. Most often, the infrared (IR) region of the electromagnetic spectrum is used, from long wavelength (thermal IR) to NIR.

The idea originates from informal experiments presented in [56]. The paper presents examples of face images of individuals, taken in the long wavelength infrared region of the spectrum ($8 - 15\mu m$), also known as thermal infrared region. The images represent the thermal emissions naturally coming from the human body. Depending on the spoofing attack material, these thermal emissions can be significantly reduced if an individual holds the spoofing attack in front of the face.

Operating in the NIR spectrum, [54] suggests that there is an apparent difference between the reflectance property of the human skin and other materials. [31] analyzes the reflectance properties of skin and artificial materials at two wavelengths: one in NIR and one in VIS spectrum. The two obtained measurements form a feature vector for a multispectral based spoofing detection. Trying to overcome the requirement for a particular distance from the sensor in, [31], [80] finds the most suitable wavelengths and trains the system with data taken at multiple distances. In [73], the authors use multispectral filters to obtain an image which presents the different radiometric response of different parts of the face under a full-spectrum active light. The distinction between real accesses and spoofing attacks is made by analyzing the gradient of the image.

Going back to the visible spectrum, [3] measures the reflectance of the skin using a high resolution, high accuracy spectrograph. Using 8 different wavelengths in the visible spectrum, [69] creates a feature vector based on the RGB values of the obtained images.

It is important to notice that in the scenarios referred to in this section, the hardware-based methods using IR sensors are used to protect face recognition system in the visible spectrum. However, these methods are even more suitable to operate alongside face recognition systems in the IR spectrum. IR and multispectral face recognition systems will be covered in Section 2.2.

Another example of a hardware-based method is the recent approach [59] which uses, the newly developed light-field camera that records the direction and intensity of each light ray. This camera renders multiple focus images in a single shot. Using this technology, it is possible to distinguish between the multiple focus levels to distinguish between 2D spoofing attacks and real faces.

The need of an additional hardware renders the hardware-based method more expensive and less convenient from deployment perspective. This requirement implies that some of them can not be used in certain applications, like for example mobile systems.

Fusion

The main motivation behind approaches proposing fusion of anti-spoofing methods, is the fact that different types of spoofing attacks have different properties and it is difficult to address all of them only with a single feature type or method. In addition, [22] has made a proof of concept that the anti-spoofing systems are unable to generalize well on unseen spoofing attacks. The discussion in the previous sections, where we state which spoofing attacks are most likely to be detected by the various categories of methods, is an argument towards this direction. Hence the trend of fusing several different anti-spoofing methods to obtain a more general counter-measure effective against a multitude of attack types.

The first attempts of fusing have been performed by [67], where the authors develop a fusion scheme at a frame and video level and apply it to a set of visual appearance cues, and [64], where the fusion of visual appearance cues is done at feature level. The authors in [75] for the first time bring the intuition that the fusion can have a bigger impact if done with complementary counter-measures, i.e. those that address different types of cues at the spoofing attacks. In the particular case, although subject to some prerequisites of the videos, motion analysis method is fused with a visual appearance method.

To measure the level of independence of two anti-spoofing systems, and thus to get an measurement of the effectiveness of their fusion, [22] proposes employing a statistical analysis based on [41]. For the same purpose, [34] proposes to count the common error rates. [34] further shows that fusing several simple anti-spoofing methods which do not involve complex inefficient classifiers may be favorable with respect to a single one which is memory and time requiring.

The trend of fusing multiple complementary anti-spoofing methods continued with [14]. While fusion at score level is the most dominant approach, future efforts should analyze what is the most effective fusion strategy, both in terms of error rates, as well as flexibility of incorporating a newly developed counter-measure into the fused system.

2.1.4 Discussion

2D spoofing attacks in VIS have attracted a lot of interest among researchers in the past years. This resulted in a large set of counter-measures belonging to different categories, with different efficiency and targeting different types of attacks. Besides this, the counter-measures differ in other important properties, like their intrusiveness and the type of input they require. We believe that summarizing the available methods based on their properties is much more important than comparing their performance, because each one is tested and works on different conditions. For this purpose, we systematized them in Table 4, grouping them by category and listing their main properties. In this way, a user can decide which method to use based on the expected spoofing attacks, types of input the system provides, as well as ease of implementation and convenience of use.

From the results published in the literature so far, we can deduce two main conclusions which may serve to direct future research.

1. Many publications have already achieved close to zero or zero error rates in spoofing detection for the three main publicly available face spoofing databases. The community has recognized the limitations of the currently existing databases, ranging from small number of identities, to small set of spoofing attacks types, to various types of biases. More challenging databases need to be created in future. Considering different materials to produce the spoofing attacks, using better quality equipment, creating more diverse illumination conditions or recording more clients are some of the ways to add to the adversity of the spoofing databases.
2. Several publications have shown that the proposed anti-spoofing methods do not generalize well on new spoofing attacks not seen during training time [22, 76]. However, good generalization capabilities should be a key feature for the anti-spoofing methods, as new types of spoofing attacks can never be anticipated. Therefore, future research effort should put an emphasis on methods which generalize well over multitude of different types of spoofing attacks.

2.2 NIR and multispectral face recognition systems

The face recognition systems which work in the infrared part of the spectrum have one major advantage over their counterparts in VIS spectrum: they are usually invariant to illumination changes in the environment. NIR face recognition systems need an active NIR light to illuminate the subject and capture the reflection of the face under that light. Examples of the robustness of these systems has been demonstrated in [84, 44]. Multispectral systems are created by fusing face recognition systems which work in different part of the spectrum, like NIR and thermal, or NIR and VIS [65]. However, the robustness to spoofing attacks of these systems has been addressed very sparsely.

The objective of this section is to study several examples of face recognition systems working in NIR, and to evaluate their vulnerability to a basic type of spoofing

attacks. We present a new publicly available multispectral face spoofing database, containing face images in NIR and VIS spectrum. The systems are evaluated when working in VIS and NIR spectrum, as well as in multispectral scenario, by fusing the scores of the VIS and NIR systems.

2.2.1 Types of attacks and databases

The attempts to spoof face recognition system in NIR spectrum are by far less numerous than similar attempts in VIS spectrum. The work in [54] presents a way to use NIR technology to detect spoofing attacks for VIS spectrum face recognition systems. Some of them even present an empirical study on the success in detecting spoofing attacks. However, none of these studies creates and evaluates spoofing attacks designated to NIR and/or multispectral face recognition systems. That is, in fact, a basic preliminary step before developing a counter-measure.

To the best of our knowledge, only [78] has studied the effect of spoofing attacks on NIR and multispectral face recognition system. The authors develop a database with 100 clients, taking simultaneously VIS spectrum and NIR images at each shot. Then, spoofing attacks are created from part of the recorded images in the two spectra, by printing them on a coarse paper. In this way, both VIS and NIR spoofing attacks are created. A disadvantage of the study on [78] is that the database is not publicly available.

To alleviate this issue, we created a new publicly available database, called Multispectral-Spoof³. The total number of clients in the database is 21. The database is recorded using a uEye camera with CMOS monochrome sensor and a resolution of 1280x1024. The images in NIR were recorded using a NIR illuminator and a NIR cut filter of 800nm attached to the camera. The images were taken in 7 different conditions: one in an uncontrolled hallway environment and 6 in office environment with natural light, ambient light, no light, illuminator spotlight from the left and from the right and 2 illuminator spotlights. 5 images in VIS and 5 images in NIR spectrum were taken under each of these conditions.

Bearing in mind that the attacker may have an access to the best quality real access samples of the clients, we selected the 3 best images from the VIS and NIR samples of each client and printed them in black and white on a normal A4 paper, using a printer with 600 dpi. Then, using the same settings as before, we recorded the printed spoofing attacks in both VIS and NIR spectrum in 3 different lighting conditions in an office environment: natural light, ambient light and 2 illuminator spotlights. For an unbiased evaluation, the clients in the database are divided into 3 non-overlapping sets for training (9 clients), development (6 clients) and testing (6 clients) purposes. Figure 5 and Figure 6 illustrate examples of real access and attack samples taken in VIS and NIR, respectively.

³ The link to download the database, together with manual face annotations, will be available as soon as this book chapter is accepted for publication.




Fig. 5 Real and spoofing attack samples from the database recorded in VIS




Fig. 6 Real and spoofing attack samples from the database recorded in NIR

2.2.2 Assessing the vulnerability

In this section, we study the effectiveness of VIS and NIR spoofing attacks in defeating VIS and NIR recognition systems. We would like to inspect whether it is possible to spoof VIS systems using NIR attacks and vice versa. First insight into this problem has been reported by [78]. The studied face recognition system [44] is based on Gabor wavelets. The authors conclude that while VIS system is vulnerable to VIS attacks and NIR system is vulnerable to NIR attacks, there are little chances that VIS attacks can bypass a NIR system and vice-versa.

We perform similar analysis using the publicly available Multispectral-Spoof database. We analyze the same recognition systems described in Section 2.1.2: GMM, LGBPHS, GJet and ISV, this time operating in two domains: VIS and NIR⁴.

The Multispectral-Spoof database contains a total of 1680 real access images (840 in VIS and 840 in NIR), as well as 3024 spoofing attack images (756 VIS and 756 NIR attacks for each of the two systems). To allow for training and evaluation of face recognition systems and following the example of Replay-Attack, 10 of the images of each client are reserved for enrollment purposes. During the evaluation,

⁴ The link to fully reproduce the results obtained here will be available as soon as this book chapter is accepted for publication.

the vulnerability of each of the systems (VIS and NIR) when exposed to the two types of attacks (VIS and NIR) was assessed.

Independent VIS and NIR systems

We firstly analyze the verification performance and the vulnerabilities of GMM-based system working in VIS mode. The score distributions for this system are given in Fig. 7, and the good separation between the distribution of the real accesses and spoofing attacks indicates that the system behaves relatively well in verification. However, Fig. 7a shows that the system is highly vulnerable to spoofing attacks recorded in VIS. More surprisingly, Fig. 7b shows that the system can be spoofed even with spoofing attacks taken in NIR spectrum, with probability of 30.56%.

Fig. 8 demonstrates similar analysis when the GMM-based system works in NIR mode. Again, the system shows relatively good verification performance. In this case the system shows low vulnerability to VIS attacks, amounting to 13.96%. The vulnerability to NIR attacks, however, goes as high as 71.8%.




Fig. 7 Score distribution of GMM based face recognition system for the samples in Multispectral-Spoof: VIS mode. Real accesses: ■; zero-effort impostors: ■; spoofing attacks: ■

Table 2 presents the verification results and the vulnerabilities for the rest of the studied face recognition systems. All of them are moderately to highly vulnerable to spoofing attacks recorded in the spectrum that they operate in. For example, SFAR for VIS systems to VIS attacks ranges from 59.26% to 74.07%. In NIR mode, the systems are even more vulnerable to NIR attacks: SFAR ranges from 71.76% to 88.89%. As can be expected, the vulnerability to attacks recorded in the other spectrum than the one the systems work in, is much lower. However, it still amounts to a considerable SFAR, especially in the case of VIS system: the SFAR for NIR attacks is between 27.78% to 38.89%. Among the studied systems, GJet appears to be the most vulnerable, while ISV shows the greatest robustness to spoofing attacks, both in VIS and NIR mode.

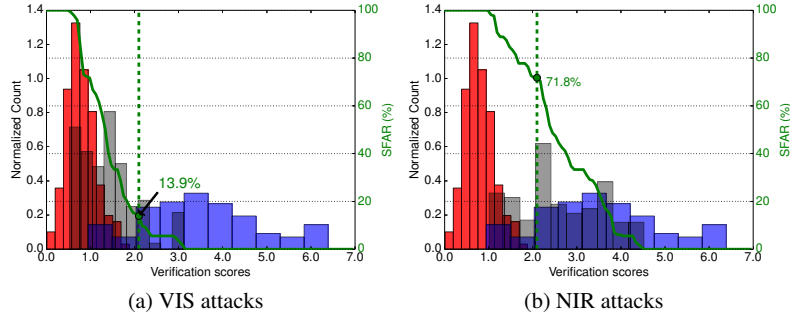


Fig. 8 Score distribution of GMM based face recognition system for the samples in Multispectral-Spoof: NIR mode. Real accesses: ■; zero-effort impostors: ■; spoofing attacks: ■

Table 2 Verification error rates and spoofing vulnerability of baseline face verification systems (in %)

system	VIS system				NIR system			
	FAR	FRR	SFAR		FAR	FRR	SFAR	
			VIS attack	NIR attack			VIS attack	NIR attack
GMM	0.78	15	62.04	30.56	0	13.96	13.89	71.76
LGBPHS	13.11	3.33	69.44	54.17	4.13	11.17	25.93	74.07
GJet	9.89	6.11	74.07	38.89	3.35	6.15	27.78	88.89
ISV	1.44	16.67	59.26	27.78	0	12.29	14.81	72.22

Multispectral system

The analysis presented in [78] is extended to a multispectral system by fusing the scores of the attacks on the two systems. If simple SUM rule is used for the score fusion, the multispectral system appears to be vulnerable to any of the two types of spoofing attacks.

In our case, we investigate three different strategies to fuse the scores of VIS and NIR systems: SUM of scores, Linear Logistic Regression (LLR) and Polynomial Logistic Regression (PLR). The vulnerabilities of the GMM-based system working in multispectral mode are given in Table 3.

Table 3 Verification error rates and spoofing vulnerability of multispectral GMM-based system (in %)

fusion method	FAR	FRR	SFAR	
			VIS attack	NIR attack
SUM	0	11.17	11.11	33.02
LLR	0	14.53	9.26	25.12
PLR	0	10.06	9.72	53.95

The results show that the vulnerability of the multispectral system is highly reduced, especially to VIS spoofing attacks. The vulnerability to NIR spoofing attacks is reduced to a lesser extent. However, the obtained SFAR has moderately high level and suggests that VIS and NIR spoofing attacks present a considerable security threat even for multispectral systems. The results for the other face recognition systems (LGBPHS, GJet and ISV) bring to similar conclusions.

2.2.3 Discussion

Research on spoofing and anti-spoofing for NIR and multispectral face recognition system is still in its infancy. We contribute to the attempts to spoof such systems by creating a publicly available VIS and NIR face spoofing database, that can be used in a multispectral setting as well. From our initial experiments we see that it is possible to spoof VIS and NIR systems with both VIS and NIR spoofing attacks. We envision three main direction for future research.

1. Multispectral-spoof database offers just the most basic spoofing attacks with printed photographs. More challenging spoofing attacks need to be created and evaluated, like 3D attacks, or image-level fusion of VIS and NIR images.
2. Multispectral systems appear to be more robust, but still not highly secure under NIR spoofing attacks. Examining different fusion strategies at different levels, fine tuning the training of the systems, fine-tuning the operating frequencies of the NIR and VIS systems and including spoofing attacks to train the fusion systems are some of the possible ways to improve the multispectral systems.
3. The set of spoofing counter-measures for these systems is very sparse. Several of the hardware-based anti-spoofing methods described in Section 2.1.3 could be readily employed for detecting spoofing attacks in NIR spectrum as well. Yet, they may still be classified as requiring additional hardware, because they operate at different wavelengths than the wavelengths used by Multispectral-Spoof database. In practice, only [78] has developed a fully software-based counter-measures for printed attacks to NIR and multispectral systems, but its efficiency to other databases and more challenging spoofing attacks is still to be tested.

Table 4 Categorization of anti-spoofing methods and overview of their main properties

Category	Method	Other category	Tested on public data	Source code	Intrusive	Type of input	Targeted attacks
Liveness detection	[51]	-	no	no	somewhat	video	print
	[71]	-	no	no	yes	video	print
	[32]	Motion analysis / Fusion	no	no	somewhat	video	print
	[8]	Visual appearance and motion analysis	Replay-Attack, CASIA-FASD	no	no	video	all attacks
Motion analysis	[33]	-	no	no	no	video	print
	[7]	-	no	no	no	video	print / warped print
	[5]	-	Print-Attack	yes	no	video	scenic print
	[4]	-	Photo-Attack	yes	no	video	scenic print / photo
Visual Appearance	[43]	-	no	no	no	video	print
	[66]	-	NUAA	no	no	image	print
	[83]	-	CASIA-FASD	no	no	image / video	print / deformed print / video
	[6]	-	no	no	no	image	print
	[26]	-	no	no	no	image	print
	[67]	Fusion	Print-Attack	no	no	image / video	print
	[64]	-	NUAA, Print-Attack	no	no	image	print
	[25]	-	Replay-Attack	no	no	image	print / photo / video
	[45]	-	NUAA	no	no	image	print
	[46]	Fusion	NUAA, Yale recaptured, Print-Attack	no	no	image	print
	[12]	-	Replay-Attack, NUAA, CASIA-FASD	yes	no	image	print, photo, video
	[77]	-	NUAA, Print-Attack, CASIA-FASD	yes	no	image	print, video
	[55]	-	no	no	no	video	print / video
	[23]	-	Replay-Attack, CASIA-FASD	yes	no	video	print, photo, video
	[36]	-	-	no	no	image	mask
[38]	-	Private database	no	no	image	mask	
[20]	-	3DMAD	yes	no	image	mask	
[57]	Fusion	3DMAD	no	no	image	mask	
Contextual Information	[52]	Liveness detection / Fusion	no	no	somewhat	video	scenic print / deformed print
	[35]	-	yes	no	no	image / video	attacks with visible medium borders
3D information	[72]	Challenge response	no	no	yes	image sequence / video	2D-attacks
	[16]	Challenge response	no	no	yes	video	2D attacks
	[30]	-	no	no	no	2 images	2D attacks
	[36]	-	-	no	no	depth image	mask
	[20]	-	3DMAD	yes	no	depth image	mask
[57]	Fusion	3DMAD	no	no	depth image	mask	
Challenge-response	[24]	Motion analysis	no	no	yes	-	-
	[2]	Motion analysis	no	no	yes	video	print
Feature learning	[76]	-	Replay-Attack, CASIA-FASD	no	no	image	all
	[48]	-	Replay-Attack, 3DMAD	no	no	image	all
Fusion	[75]	Motion Analysis / Visual Appearance	Print-Attack	no	no	video	print
	[34]	-	Replay-Attack	yes	no	video	scenic attacks

Face recognition system

3 Face recognition systems under 3D spoofing attacks

It is repeatedly stated in the previous sections that an attacker can attempt to gain access through a 2D face recognition system (visible, near-infrared or multispectral) simply by using printed photos or recorded videos of valid users. It is also reported that most of these attack types devised until today can be successfully averted by using various anti-spoofing methods.

A substantial part of the work on spoofing detection capabilities for face is based on the flatness of the surface in front of the sensor during an attack. For instance, the motion analysis techniques detailed in Section 2.1.3 rely on the assumption of shape difference between an actual face and a spoofing attack instrument such as a paper or a tablet computer in order to distinguish motion patterns of a real person from an attacker. In a similar fashion, 3D shape information either extracted from multiple-view images or acquired using a 3D sensor (Section 2.1.3) can be exploited to positively detect 2D attacks. For instance, in [68], 3D data captured with a low-cost sensor is utilized to locate the face in an image as well as to test its authenticity.

These types of methods that rely on the assumption of a planar surface that displays a face image in front of the sensor are ineffective in case of 3D facial mask attacks [18]. Although the advancements in 3D acquisition and manufacturing technologies make this kind of attacks as untroublesome as their 2D counterparts, there have not been many studies published addressing this issue. In this section, an overview of the existing work is presented for several kinds of 3D attacks, face recognition systems and spoofing counter-measures.

3.1 Types of attacks and databases

The earliest research works that target 3D attacks only aim to distinguish between facial skin and mask materials without analyzing the spoofing performances of the masks because they approach this problem as in an evasion or disguise scenario [56, 54]. The masks utilized for the experiments are not necessarily replicas of valid users.

Claiming that fake, by its definition, is indistinguishable for human eyes and visible spectrum cannot be sufficient to detect the attacks, a small group of studies follow the footsteps of early pioneers and propose multispectral analyses [31, 81] for mask and real face classification. The experiments in [31] are done on directly mask materials. In [81], some face-like masks are produced but they do not mimic any real person. Unfortunately, no public database has been made available for further investigation.

Recently, another line of research in 3D spoofing has emerged for which the attacks are realized with 3D printed masks of valid users. Firstly, Kose et al. published a series of studies [36, 37, 39, 38] on 3D mask attacks for which a non-public database of 16 users is utilized. In order to construct this database that is called Morpho database, a 3D face model of each client is captured by a 3D laser scanner. It




Fig. 9 (a) Example shots from Morpho: The top row shows a real access from a user in gray-scale texture (2D), depth map (2.5D) and 3D model format while an attacker wearing the same users mask is displayed in the bottom (b) Example paper-craft mask from 3DMAD (c) 17 Wearable resin masks from 3DMAD [20]

consists of 207 real access and 199 mask attack samples as both 2D images and 3D scans (Fig. 9a).

Morpho database did certainly bring on a significant breakthrough and momentum in 3D spoofing attack research. Still, it was lacking a very crucial characteristic that is publicness. Taking this shortcoming into account, Erdogmus et al. collected the first public spoofing database with facial masks, called 3D Mask Attack Database (3DMAD) [19] and published a couple of spoofing and counter-measure analyses on several face recognition systems [19, 20]. The database contains 76500 real access and mask attack frames of 17 users, recorded using Microsoft Kinect.

The masks used for Morpho database were printed using 3D laser scans of valid users. The acquisition process with such scanners requires cooperation since it is very sensitive movement and has range limitations. This makes the attack scenario less realistic. On the other hand, the masks for 3DMAD are manufactured using only 2D images of users via a private company which is specialized in facial reconstruction and in transforming 2D portraiture into 3D sculptures. Using this service, it is possible to construct a 3D face model from frontal and profile images of a person which can be easily obtained from a distance or found on the Internet. Once the 3D models are constructed, they can be turned into masks of various sizes and materials.

For 3DMAD, a life-size wearable mask and a paper-craft mask is manufactured for each user (Fig. 9b and 9c). Paper-craft masks can be just printed out and hand-crafted, so they are not recorded but made available within the database for the use of the biometrics community. Using Microsoft Kinect for Xbox 360, videos are recorded for real accesses and attacks with wearable hard resin masks. Since Kinect can capture both color and depth data, the database consists of both RGB and depth samples. Because of this, it enables researchers to analyze the vulnerability of 3D face recognition systems to mask attacks and to devise countermeasures in 3D.

The two above-mentioned databases constitute the backbone of research on 3D spoofing attacks that investigate the ability of masks to spoof face recognition sys-




Fig. 10 Score distributions of genuine and impostor scores on the development set and mask attack scores on the test set of 3DMAD using (a) ISV [19] (b) SRC [57], for 2D face verification

tems and the possible anti-spoofing techniques which will be detailed in the following subsections.

3.2 System vulnerabilities

With both Morpho and 3DMAD databases, vulnerabilities against spoofing with 3D masks have been analyzed extensively for 2D, 2.5D and 3D face recognition systems.

In [39], a 2D system based on Local Binary Pattern (LBP) and a 3D system based on Thin Plate Spline (TPS) warping are analyzed for their robustness against mask attacks using the Morpho database. While both system performances decline remarkably as the attacks are introduced, 3D face recognition system which is completely based on 3D facial shape analysis is found to be affected more (EER increases from 1.8% to 25.1%) than the 2D system (EER increases from 4.7% to 9.3%). This is an expected outcome since the masks in Morpho database are highly precise in shape but have only gray-scale texture. These findings are revised and extended in [37] with the addition of an LBP based 2.5D face recognition system for which the EER increase from 7.27% in normal mode to 14.26% under spoofing attacks.

Similarly, 3DMAD is also assessed with regards to its spoofing ability on various face recognition systems. Firstly in [19], an Inter Session Variability (ISV) based 2D face recognition algorithm is tried and 65.7% of the mask attacks are found to be successful at EER threshold calculated on the development set of the database. The FAR at the same threshold would increase from 1.06% to 13.99% if mask attacks are included in the probe partition together with the zero-effort impostors. The score distribution of the real access, zero-effort and mask attack impostors are given in Fig. 10a. The authors extend their study in [20] to include an ISV based

2.5D and an Iterative Closest Point (ICP) based 3D face recognition systems as well as all three baseline systems in [37]. Furthermore, spoofing performances are measured and reported separately for each mask. The experimental results reveal that the spoofing performances differ greatly not only between masks but also between modes and algorithms. Additionally, it is observed that the vulnerability to mask attacks is greater for more successful face verification algorithms that can generalize well to variations in facial appearance.

In a more recent work [57], 3DMAD masks are tested against another 2D face recognition algorithm which is based on the Sparse Representation Classifier (SRC) and 84.12% of the masks are found to be able to access the system at EER threshold (Fig. 10b).

All these findings expose that 3D mask attacks can be a real threat to all types of face recognition systems in 2D, 2.5D or 3D and serious measures should be taken in order to detect and prevent them.

3.3 Spoofing counter-measures

Several methods have been proposed to detect 3D mask attacks in both 2D and 2.5D, mainly focusing on differences between micro-texture properties of mask materials and facial skin.

In [36], Kose et al. report 88.1% and 86.0% accuracies on Morpho database with texture images (2D) and depth maps (2.5D), respectively, by concatenating histograms of different types of LBP and classifying them with an SVM classifier. Later in [39], they also try to fuse the two modes (image and depth map) at both feature and score level and reach 93.5% accuracy. Other than micro-texture analysis via LBP, they also experiment with reflectance analysis to detect 3D mask attacks in [38] and report 94.47% classification success. Finally, by fusing micro-texture and reflectance analyses in both 2D and 2.5D, an accuracy of 98.99% is reached [40].

Spoofing counter-measure studies with 3DMAD database also mainly revolves around LBP-based classification algorithms. In [19], the effectiveness of LBP based features extracted from color and depth images to detect the mask attacks is analyzed. The results suggest that LBP features extracted from overlapping blocks give better results which achieve HTER of 0.95% and 1.27% with images and depth maps separately. This study is elaborated further in [20] with best performance obtained by regular block-based LBP and a Linear Discriminant Analysis (LDA) classifier at $0.12 \mp 0.47\%$ and $3.91 \mp 6.04\%$ HTER for 2D and 2.5D.

In addition to LBP, Raghavendra et al. propose to utilize Binarized Statistical Image Features (BSIF) to capture prominent micro-texture features [57] in 2D images both for the whole face (global) and the periocular (local) region. The LBP and BSIF features for each region are classified and the final scores are fused by weighted voting which results in an HTER of 4.78%. Later in [58], the same pro-

toocol is also applied for 2.5D and the findings are incorporated via weighted score fusion. This addition is reported to push the HTER down to 0.03%.

3.4 Discussion

Utilization of 3D masks for face spoofing have certainly become easier and cheaper. Many recent studies mentioned above have revealed the vulnerability of 2D, 2.5D and 3D face recognition systems to such attacks. Additionally, many countermeasures have been proposed. However, as shown in [20], even though they are manufactured in similar ways, masks can behave very differently in various settings, making it very difficult to find one single solution that works for all.

Furthermore, in each of currently existing work, mask attack samples are utilized for training the anti-spoofing systems. This is not a realistic assumption for a biometric system since it cannot employ a different anti-spoofing module for each different mask. Worse still, it is always possible to encounter new and unseen types of masks.

The anti-spoofing methods targeting 3D masks have been added to Table 4, together with the anti-spoofing methods for 2D attacks described in Section 2.1.3. Table 4 thus represents a comprehensive summarization of all the efforts in face spoofing detection in visible spectrum that have been published so far.

4 Conclusions

Spoofing attacks are one of the most important reasons why face recognition may have a limited application in conditions where supervision is not possible. Face spoofing attacks have been proved to be effective for face recognition systems in VIS spectrum in many occasions, including several face spoofing databases. So far, many counter-measures have been developed, and each of them tackles the problem from a different perspective. As a result, most of these counter-measures are effective just for a subset of the spoofing attack types. Having in mind the limitation of the currently available databases, as well as the possibility of new spoofing attacks appearing in future, more research efforts are needed to enhance the generalization capabilities of the counter-measures.

The work in spoofing face recognition systems in NIR is not as extensive. However, the newly developed Multispectral-Spoof database which includes VIS and NIR attacks, demonstrates the vulnerability of both VIS and NIR systems to such attacks. Employing these systems in multispectral scenario significantly reduces the risks. Yet, development of suitable counter-measures is needed to provide acceptable security levels for multispectral face recognition systems.

The published research in anti-spoofing for face recognition rarely comes with data or source code that can be reproduced. This poses difficulties when comparing

the performance of counter-measures. We would like emphasize the importance of publishing fully reproducible spoofing databases and counter-measures, as this will be of great benefit for building upon existing solutions and development to encourage the practice of new ones. In this chapter, we explicitly pointed out to solutions which are fully reproducible and we would like to encourage this practice for the future work.

References

1. Akhtar, Z., Fumera, G., Marcialis, G.L., Roli, F.: Evaluation of serial and parallel multibiometric systems under spoofing attacks. In: 5th IEEE International Conference on Biometrics: Theory, Applications and Systems (2012)
2. Ali, A., Deravi, F., Hoque, S.: Spoofing attempt detection using gaze colocation. In: Biometrics Special Interest Group (BIOSIG), 2013 International Conference of the (2013)
3. Angelopoulou, E.: Understanding the color of human skin. pp. 243–251 (2001)
4. Anjos, A., Chakka, M.M., Marcel, S.: Motion-based counter-measures to photo attacks in face recognition. *Institution of Engineering and Technology Journal on Biometrics* (2013)
5. Anjos, A., Marcel, S.: Counter-measures to photo attacks in face recognition: a public database and a baseline. In: International Joint Conference on Biometrics 2011 (2011)
6. Bai, J., et al.: Is physics-based liveness detection truly possible with a single image? In: IEEE International Symposium on Circuits and Systems (ISCAS) (2010)
7. Bao, W., Li, H., Li, N., Jiang, W.: A liveness detection method for face recognition based on optical flow field. 2009 International Conference on Image Analysis and Signal Processing pp. 223–236 (2009)
8. Bharadwaj, S., Dhamecha, T., Vatsa, M., Singh, R.: Computationally efficient face spoofing detection with motion magnification. In: Computer Vision and Pattern Recognition Workshops (CVPRW), 2013 IEEE Conference on, pp. 105–110 (2013)
9. Cardinaux, F., Sanderson, C., Marcel, S.: Comparison of mlp and gmm classifiers for face verification on xm2vts. In: Proceedings of the 4th International Conference on AVBPA. University of Surrey, Guildford, UK (2003)
10. Chetty, G., Wagner, M.: Audio-visual multimodal fusion for biometric person authentication and liveness verification. In: Proceedings of the 2005 NICTA-HCSNet Multimodal User Interaction Workshop - Volume 57, pp. 17–24. Australian Computer Society, Inc. (2006)
11. Chetty, G., Wagner, M.: Multi-level liveness verification for face-voice biometric authentication. In: In Biometrics symposium 2006 (2006)
12. Chingovska, I., Anjos, A., Marcel, S.: On the effectiveness of local binary patterns in face anti-spoofing. In: Proceedings of the 11th International Conference of the Biometrics Special Interest Group (2012)
13. Chingovska, I., Anjos, A., Marcel, S.: Biometrics evaluation under spoofing attacks. *IEEE Transactions on Information Forensics and Security* **9**(12), 2264–2276 (2014)
14. Chingovska, I., et al.: The 2nd competition on counter measures to 2d face spoofing attacks. In: International Conference of Biometrics (ICB) (2013)
15. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on, vol. 1, pp. 886–893 vol. 1 (2005)
16. De Marsico, M., Nappi, M., Riccio, D., Dugelay, J.: Moving face spoofing detection via 3d projective invariants. In: Biometrics (ICB), 2012 5th IAPR International Conference on (2012)
17. Duc, N.M., Minh, B.Q.: Your face is not your password. Black Hat Conference (2009)
18. Erdogmus, N., Marcel, S.: Spoofing 2d face recognition systems with 3d masks. In: International Conference of the Biometrics Special Interest Group (BIOSIG), pp. 1–8 (2013)

19. Erdogmus, N., Marcel, S.: Spoofing in 2d face recognition with 3d masks and anti-spoofing with kinect. In: International Conference on Biometrics: Theory, Applications and Systems (BTAS), pp. 1–6. IEEE (2013)
20. Erdogmus, N., Marcel, S.: Spoofing face recognition with 3d masks. IEEE Transactions on Information Forensics and Security **9**(7), 1084–1097 (2014)
21. Flynn, P.J., Faltemier, T., Bowyer, K.W.: 3d face recognition. In: A.K. Jain, P. Flynn, A.A. Ross (eds.) Handbook of Biometrics, pp. 293–313 (2008)
22. de Freitas Pereira, T., Anjos, A., De Martino, J.M., Marcel, S.: Can face anti-spoofing countermeasures work in a real world scenario? In: International Conference on Biometrics (2013)
23. de Freitas Pereira, T., et al.: Face liveness detection using dynamic texture. EURASIP Journal on Image and Video Processing **2014**:2 (2014)
24. Frischholz, R., Werner, A.: Avoiding replay-attacks in a face recognition system using head-pose estimation. In: Analysis and Modeling of Faces and Gestures, 2003. AMFG 2003. IEEE International Workshop on (2003)
25. Galbally, J., Marcel, S., Fierrez, J.: Image quality assessment for fake biometric detection: Application to iris, fingerprint and face recognition. IEEE Trans. on Image Processing **23**(2), 710–724 (2014)
26. Gao, X., Ng, T.T., Bo, Q., Chang, S.F.: Single-view recaptured image detection based on physics-based features. In: IEEE International Conference on Multimedia & Expo (ICME) (2010)
27. Günther, M., Haufe, D., Würtz, R.P.: Face recognition with disparity corrected Gabor phase differences. In: Artificial Neural Networks and Machine Learning, *Lecture Notes in Computer Science*, vol. 7552, pp. 411–418. Springer Berlin (2012)
28. Jain, A.K., Ross, A.: Handbook of Biometrics, chap. Introduction to Biometrics. Springer-Verlag (2008)
29. Johnson, P.A., Tan, B., Schuckers, S.: Multimodal fusion vulnerability to non-zero (spoof) imposters. In: IEEE International Workshop on Information Forensics and Security (2010)
30. Kim, S., Yu, S., Kim, K., Ban, Y., Lee, S.: Face liveness detection using variable focusing. In: Biometrics (ICB), 2013 International Conference on (2013)
31. Kim, Y., Na, J., Yoon, S., Yi, J.: Masked fake face detection using radiance measurements. Journal of the Optical Society of America A **26**(4), 760–766 (2009)
32. Kollreider, K., Fronthaler, H., Bigun, J.: Verifying liveness by multiple experts in face biometrics. In: Computer Vision and Pattern Recognition Workshops, 2008. CVPRW '08. IEEE Computer Society Conference on, pp. 1–6 (2008)
33. Kollreider, K., Fronthaler, H., Bigun, J.: Non-intrusive liveness detection by face images. Image and Vision Computing **27**(3), 233–244 (2009)
34. Komulainen, J., Anjos, A., Hadid, A., Pietikainen, M., Marcel, S.: Complementary countermeasures for detecting scenic face spoofing attacks (2013)
35. Komulainen, J., Hadid, A., Pietikainen, M.: Context based face anti-spoofing. In: Biometrics: Theory, Applications and Systems (BTAS), 2013 IEEE Sixth International Conference on, pp. 1–8 (2013)
36. Kose, N., Dugelay, J.L.: Countermeasure for the protection of face recognition systems against mask attacks. In: International Conference and Workshops on Automatic Face and Gesture Recognition (FG), pp. 1–6. IEEE (2013)
37. Kose, N., Dugelay, J.L.: On the vulnerability of face recognition systems to spoofing mask attacks. In: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 2357–2361. IEEE (2013)
38. Kose, N., Dugelay, J.L.: Reflectance analysis based countermeasure technique to detect face mask attacks. In: International Conference on Digital Signal Processing (DSP), pp. 1–6. IEEE (2013)
39. Kose, N., Dugelay, J.L.: Shape and texture based countermeasure to protect face recognition systems against mask attacks. In: IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 111–116. IEEE (2013)
40. Kose, N., Dugelay, J.L.: Mask spoofing in face recognition and countermeasures. Image and Vision Computing **32**(10), 779–789 (2014)

41. Kuncheva, L.I., Whitaker, C.J.: Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy. *Machine Learning* **51**(2), 181–207 (2003)
42. L. Thalheim J. Krissler, P.M.Z.: Body check: Biometric access protection devices and their programs put to test. *Heise Online* (2002)
43. Li, J., et al.: Live face detection based on the analysis of fourier spectra. *Biometric Technology for Human Identification* (2004)
44. Li, S.Z., Chu, R., Liao, S., Zhang, L.: Illumination invariant face recognition using near-infrared images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **29**(4), 627–639 (2007)
45. Määttä, J., Hadid, A., Pietikäinen, M.: Face spoofing detection from single images using micro-texture analysis. In: *International Joint Conference on Biometrics*, pp. 1–7 (2011)
46. Määttä, J., Hadid, A., Pietikäinen, M.: Face spoofing detection from single images using texture and local shape analysis. *IET Biometrics* **1**, 3–10 (2012)
47. Mansfield, A.J., Wayman, J.L.: Best practices in testing and reporting performance (2002)
48. Menotti, D., Chiachia, G., Pinto, A., Schwartz, W., Pedrini, H., Falcao, A., Rocha, A.: Deep representations for iris, face, and fingerprint spoofing detection. *Information Forensics and Security, IEEE Transactions on* (99), 1–1 (2015)
49. Ojala, T., Pietikainen, M., Maenpaa, T.: Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **24**(7), 971–987 (2002)
50. Oren, M., Nayar, S.K.: Generalization of the lambertian model and implications for machine vision. *International Journal of Computer Vision* **14**(3), 227–251 (1995)
51. Pan, G., Sun, L., Wu, Z., Lao, S.: Eyeblick-based anti-spoofing in face recognition from a generic webcam. In: *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pp. 1–8 (2007)
52. Pan, G., Sun, L., Wu, Z., Wang, Y.: Monocular camera-based face liveness detection by combining eyeblink and scene context. *Telecommunication Systems* **47**(3-4), 215–225 (2011)
53. Parziale, G., Dittman, J., Tistarelli, M.: Analysis and evaluation of alternatives and advanced solutions for system elements. *BioSecure D 9.1.2* (2005)
54. Pavlidis, I., Symosek, P.: The imaging issue in an automatic face/disguise detection system. In: *Computer Vision Beyond the Visible Spectrum: Methods and Applications, 2000. Proceedings. IEEE Workshop on*, pp. 15–24 (2000)
55. Pinto, A.d.S., Pedrini, H., Schwartz, W.R., Rocha, A.: Video-based face spoofing detection through visual rhythm analysis. In: *25th Conference on Graphics, Patterns and Images* (2012)
56. Prokoski, F.J.: Disguise detection and identification using infrared imagery. pp. 27–31 (1983)
57. Raghavendra, R., Busch, C.: Novel presentation attack detection algorithm for face recognition system: Application to 3d face mask attack. In: *International Conference on Image Processing (ICIP)*, pp. 323–327. *IEEE* (2014)
58. Raghavendra, R., Busch, C.: Robust 2d/3d face mask presentation attack detection scheme by exploring multiple features and comparison score level fusion. In: *International Conference on Information Fusion (FUSION)*, pp. 1–7. *IEEE* (2014)
59. Raghavendra, R., Raja, K., Busch, C.: Presentation attack detection for face recognition using light field camera. *Image Processing, IEEE Transactions on* (99), 1–1 (2015)
60. Rodrigues, R., Kamat, N., Govindaraju, V.: Evaluation of biometric spoofing in a multimodal system. In: *Biometrics: Theory Applications and Systems (BTAS), 2010 Fourth IEEE International Conference on* (2010)
61. Rodrigues, R.N., Ling, L.L., Govindaraju, V.: Robustness of multimodal biometric fusion methods against spoofing attacks. *Journal of Visual Languages and Computing* **20**(3), 169–179 (2009)
62. Ross, A., Nandakumar, K., Jain, A.K.: *Handbook of Biometrics*, chap. Introduction to multi-biometrics. Springer-Verlag (2008)
63. Schuckers, S.: *Encyclopedia of Biometrics*, chap. Liveness Detection: Fingerprint, pp. 924–931. Springer-Verlag (2009)
64. Schwartz, W.R., Rocha, A., Pedrini, H.: Face Spoofing Detection through Partial Least Squares and Low-Level Descriptors. In: *International Joint Conference on Biometrics* (2011)

65. Socolinsky, D.A.: Multispectral face recognition. In: A.K. Jain, P. Flynn, A.A. Ross (eds.) *Handbook of Biometrics*, pp. 293–313 (2008)
66. Tan, X., et al.: Face liveness detection from a single image with sparse low rank bilinear discriminative model. In: *ECCV (6)*, pp. 504–517 (2010)
67. Tronci, R., et al.: Fusion of multiple clues for photo-attack detection in face recognition systems. In: *IJCB*, pp. 1–6 (2011)
68. Tsalakanidou, F., Dimitriadis, C., Malassiotis, S.: A secure and privacy friendly 2d+3d face authentication system robust under pose and illumination variation. In: *Image Analysis for Multimedia Interactive Services, 2007. WIAMIS '07. Eighth International Workshop on*, pp. 40–40 (2007)
69. Vink, J., Gritti, T., Hu, Y., de Haan, G.: Robust skin detection using multi-spectral illumination. In: *Automatic Face Gesture Recognition and Workshops (FG 2011), 2011 IEEE International Conference on*, pp. 448–455 (2011)
70. Wallace, R., McLaren, M., McCool, C., Marcel, S.: Inter-session variability modelling and joint factor analysis for face authentication. In: *International Joint Conference on Biometrics (2011)*
71. Wang, L., Ding, X., Fang, C.: Face live detection method based on physiological motion analysis. *Tsinghua Science and Technology* **14**(6), 685–690 (2009)
72. Wang, T., Yang, J., Lei, Z., Liao, S., Li, S.Z.: Face liveness detection using 3D structure recovered from a single camera. In: *Biometrics (ICB), 2013 International Conference on (2013)*
73. Wang, Y., Hao, X., Hou, Y., Guo, C.: A new multispectral method for face liveness detection. In: *Pattern Recognition (ACPR), 2013 2nd IAPR Asian Conference on*, pp. 922–926 (2013)
74. Wu, H.Y., Rubinstein, M., Shih, E., Guttag, J., Durand, F., Freeman, W.T.: Eulerian video magnification for revealing subtle changes in the world. *ACM Trans. Graph. (Proceedings SIGGRAPH 2012)* **31**(4) (2012)
75. Yan, J., Zhang, Z., Lei, Z., Yi, D., Li, S.Z.: Face liveness detection by exploring multiple scenic clues. In: *12th International Conference on Control, Automation, robotics and Vision (ICARCV 2012), China (2012)*
76. Yang, J., Lei, Z., Li, S.Z.: Learn convolutional neural network for face anti-spoofing. *CoRR abs/1408.5601* (2014)
77. Yang, J., Lei, Z., Liao, S., Li, S.: Face liveness detection with component dependent descriptor. In: *Biometrics (ICB), 2013 International Conference on*, pp. 1–6 (2013)
78. Yi, D., Lei, Z., Zhang, Z., Li, S.: Face anti-spoofing: Multi-spectral approach. In: S. Marcel, M.S. Nixon, S.Z. Li (eds.) *Handbook of Biometric Anti-Spoofing, Advances in Computer Vision and Pattern Recognition*, pp. 83–102. Springer London (2014)
79. Zhang, W., et al.: Local gabor binary pattern histogram sequence (lgbphs): A novel non-statistical model for face representation and recognition. In: *Proceedings of the Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1 - Volume 01, ICCV '05*, pp. 786–791. IEEE Computer Society (2005)
80. Zhang, Z., Yi, D., Lei, Z., Li, S.Z.: Face liveness detection by learning multispectral reflectance distributions. pp. 436–441 (2011)
81. Zhang, Z., Yi, D., Lei, Z., Li, S.Z.: Face liveness detection by learning multispectral reflectance distributions. In: *IEEE International Conference on Automatic Face & Gesture Recognition and Workshops (FG)*, pp. 436–441. IEEE (2011)
82. Zhao, G., Pietikäinen, M.: Dynamic texture recognition using local binary patterns with an application to facial expressions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **29**(6), 915–928 (2007)
83. Zhiwei, Z., et al.: A face antispoofing database with diverse attacks. In: *Proceedings of the 5th IAPR International Conference on Biometrics (ICB'12), New Delhi, India (2012)*
84. Zou, X., Kittler, J., Messer, K.: Ambient illumination variation removal by active near-ir imaging. In: D. Zhang, A.K. Jain (eds.) *Advances in Biometrics, Lecture Notes in Computer Science*, vol. 3832, pp. 19–25 (2005)