Spoofing Faces Using Makeup: An Investigative Study

Cunjian Chen
Michigan State University
cunjian@msu.edu

Antitza Dantcheva
Inria Méditerranée
Antitza.Dantcheva@inria.fr

Thomas Swearingen, Arun Ross*
Michigan State University
{swearin3, rossarun}@msu.edu

Abstract

Makeup can be used to alter the facial appearance of a person. Previous studies have established the potential of using makeup to obfuscate the identity of an individual with respect to an automated face matcher. In this work, we analyze the potential of using makeup for spoofing an identity, where an individual attempts to impersonate another person's facial appearance. In this regard, we first assemble a set of face images downloaded from the internet where individuals use facial cosmetics to impersonate celebrities. We next determine the impact of this alteration on two different face matchers. Experiments suggest that automated face matchers are vulnerable to makeup-induced spoofing and that the success of spoofing is impacted by the appearance of the impersonator's face and the target face being spoofed. Further, an identification experiment is conducted to show that the spoofed faces are successfully matched at better ranks after the application of makeup. To the best of our knowledge, this is the first work that systematically studies the impact of makeup-induced face spoofing on automated face recognition.

1. Introduction

Biometrics refers to the automated recognition of individuals based on their biological traits such as face, fingerprints and iris. A typical biometric system acquires the biometric data of an individual using a sensor; extracts a set of salient features from the data; and uses these features to determine or verify the identity of an individual [11]. In spite of its advantages, a biometric system is vulnerable to spoofing, where an adversary can spoof the biometric trait of another individual in order to circumvent the system [16, 19, 2, 25]. Unlike obfuscation, which entails deliberately obscuring one's own identity, spoofing entails taking on another person's identity, with the purpose of accessing privileges and resources associated with that person [17, 21]. Spoofing, in the context of face recognition,

Figure 1. Examples of typical spoof attacks previously considered in the literature (images obtained from [17]). An attacker presents a photograph (a) or a mask (b) to the biometric system.

can be accomplished by presenting photographs [1, 21], videos [1, 19, 21] or 2D and 3D-masks to a face recognition system [13] (as seen in Figure 1).

In this work we determine whether facial cosmetics can be used by an adversary to launch a spoof attack. Unlike spoof attacks based on photographs, videos and masks, makeup-induced spoofing can be relatively difficult to detect since makeup is widely used for cosmetic purposes. Thus, it is necessary to understand if makeup-induced spoofing can confound an automated face recognition system, *i.e.*, is the recognition performance of automated face matchers impacted by this type of spoof attack?

In order to conduct our analysis, we first assemble a dataset consisting of face images of female subjects who apply makeup to transform their appearance in order to resemble celebrities. These images are extracted from videos available on YouTube. The subjects here are **not trying to deliberately deceive** an *automated* face recognition system; rather, their intention is to impersonate a target celebrity from a *human vision* perspective.

Besides assembling the dataset, the contributions of this work include the following: (a) We define two spoofing indices to quantify the potential of using makeup for face spoofing; (b) We test the vulnerability of face recognition systems to makeup induced spoofing based on these indices; (c) We conduct an identification experiment to demonstrate the potential of spoofing a target face through the use of makeup. To the best of our knowledge, this is the first work

*Corresponding Author

to systematically study this effect.

1.1. Background and related work

Recent work has demonstrated the impact of *commonly used* facial makeup on automated face recognition systems [6, 8, 9, 12, 5], and automated face-based gender and age estimation systems [4]. Makeup can be used to alter the perceived (a) facial shape; (b) nose, mouth, eye and eye brow shape; (c) nose, mouth, eye and eye brow size; (d) facial contrast and (e) facial skin quality and color. It can also be used to conceal wrinkles, dark shadows and circles underneath the eyes, and camouflage birth moles, scars and tattoos [18].

Considering the widespread use of makeup and its implications in altering facial appearance (*e.g.*, facial aesthetics [7]), in this work we focus on the use of makeup for spoofing. Unlike previous work [6, 8, 9], where *commonly used* makeup was observed to affect face recognition systems by obscuring a person’s identity, here we consider the scenario where makeup is used by an individual to mimic the facial appearance of another individual (see Figure 2).

Current face spoof detection schemes either rely on physiological cues such as eye blinking, mouth movements, and macro- and micro-expression changes [19, 16], or textural attributes of the face image [16, 24, 23, 14]. But none of these methods represent a viable mechanism for detecting makeup induced spoofing (especially since makeup is widely used). Also, in contrast to other face alteration techniques such as plastic surgery, makeup is non-permanent and cost efficient. This makes makeup-based spoofing a realistic threat to the integrity of a face recognition system.

The rest of the paper is organized as follows. Section 2 discusses the assembled spoofing dataset. Section 3 introduces the two spoofing indices for quantifying the effect of makeup on automated face recognition. Section 4 discusses

face recognition methods used in this study to evaluate the impact of makeup induced spoofing. Section 5 presents the related experiments. Results of the experiments are discussed in Section 6, followed by a summary of the paper in Section 7.

2. Makeup Induced Face Spoofing (MIFS) Dataset

In order to investigate the problem of makeup induced face spoofing, we first assemble a dataset consisting of 107 makeup-transformations taken from random YouTube makeup video tutorials. We refer to this dataset as the Makeup Induced Face Spoofing (MIFS) dataset.¹ There are two *before-makeup* and two *after-makeup* images per subject. Since each subject is attempting to spoof a target identity, we also have two face images of the target identity from the Web. Thus, this dataset has three sets of face images: images of a subject before makeup; images of the same subject after makeup with the intention of spoofing; and images of the target subject who is being spoofed. *However, it is important to note that the target images are not necessarily those used by the spoofer as a reference during the makeup transformation process.*² This is important to point out because the spoofed celebrities can often change their facial appearance and this will have an effect on the match score between the after-makeup image of the impersonator and the target image of the celebrity. When we searched the Web for face images of the target identity, we tried to select images that most resembled the after-makeup image.

All the acquired images are subjected to face cropping. This routine eliminates hair and accessories [6]. Examples of cropped images, based on a Commercial Off-The-Shelf (COTS) face detector, are shown in Figure 3.

We make the following observations about this dataset: (a) the subjects in the before and after makeup sets do not overlap with subjects in the target set; (b) there are duplicate identities of subjects attempting to spoof — this is because there are subjects attempting to spoof different target identities (see Figure 2); (c) there are duplicate identities in the target set — this is because there are multiple subjects attempting to spoof the same target identity (see Figure 4); (d) the images in the dataset include variations in expression, illumination, pose, resolution and quality.

Makeup-transformation: In the MIFS dataset, subjects use different types of makeup to alter their appearance and resemble a target identity. While the makeup application process varies across the dataset (depending upon the sub-

Figure 2. The subject on the top attempts to resemble identities in the bottom row (labeled “Target Subjects”) through the use of makeup. The result of these attempts can be seen in the second row. Images were obtained from the WWW.

¹The MIFS dataset is available at www.antitza.com/makeup-datasets.html

²Note that the makeup video tutorials do not include images of the target identity, if any, used by the subject during the spoofing process. In fact, it is likely that some subjects are attempting to resemble the target identity from memory.

Figure 3. Examples of images in the MIFS dataset after cropping. Here, (a), (b) and (c) represent three spoofing attempts. In each case, the image on the left shows the subject before makeup, the one in the middle is the subject after makeup, and the image on the right is the target identity that the subject is attempting to impersonate (see text for explanation).

ject’s face image and the target face image), we make a few general observations here. Generally, makeup-foundation is applied on the face to create a complexion that is similar to the skin color of the target subject. Face powder is then used to fix the foundation and prevent shininess, allowing for an even, uniform appearance of the face. In the next crucial step, a contouring technique is used to mimic the key characteristic features of the target face (e.g., high cheek bones, slim face, presence of beard). Specifically, brush strokes of very dark powder (e.g., bronzer) create the effect of shadows or concave facial features (e.g., underneath the cheek bones or at the periphery of the face), while brush strokes of very bright cream (e.g., highlighter) create the illusion of convex and prominent facial features. The contours are then blended with the foundation using a brush or the fingers. Facial hair, such as beard and moustache, are usually painted with brown or black eye pencils. The mouth area is then altered to resemble that of the target face by either augmenting it (painting the area around the mouth using the target’s lip color and drawing new contours around it) or minimizing it (covering part of the lips with foundation and drawing new contours resembling the target). Similarly, the shape and size of the eye region is altered by using dark eye-pencils and white highlighter-pencils, which can either extend the eyes by painting new eye-contours around the initial ones or minimize the eyes by painting within the waterline. The periocular region is then contoured using dark

Figure 4. Examples of images in the MIFS dataset after cropping. Here, (a) and (b) represent two spoofing attempts where two different subjects (left) apply makeup (middle) to resemble the same target identity (right).

and/or bright eye-shadow or cream to capture the shape of the target’s eye (e.g., hooded or deep set eyes).

In general, such transformation requires extensive quantities of different cosmetic products compared to what is commonly used. Therefore, makeup-palettes with a variety of colors and shades are expected to be used in such makeup-transformations.

3. Face Recognition and Spoofing Metrics

To quantify the impact of makeup-based identity transformation on face recognition, we propose and define two *spoofing indices* (*SIs*). The following notations are used: the complete set of subjects in the dataset is denoted as P , the set of before-makeup images as B^P , the set of after-makeup images as A^P , and the set of target identities as T^P . We note that T^P does not include identities from A^P (and thus B^P). For each subject p , we have the following image samples: $\{B_1^p, B_2^p\} \in B^P$, $\{A_1^p, A_2^p\} \in A^P$, and $\{T_1^p, T_2^p\} \in T^P$. Let $\phi(x, y)$ denote the similarity match score between two images x and y as computed by a matcher (the greater the value, the higher the similarity between two faces). The similarity scores are normalized in the $[0, 1]$ interval.

A spoofing attack can be deemed to be successful for subject p , when the similarity score between the after-makeup images, A_i^p , and the target images, T_j^p , (i.e., $\phi(A_i^p, T_j^p)$) increases. However, it is not easy to make this assessment, since any change in match score has to be viewed with respect to the entire score distribution of the matcher (and not just the absolute change in value). Hence, we consider two spoofing indices.

The two spoofing indices that we introduce below describe the similarity score $\phi(A_i^p, T_j^p)$ with respect to two types of genuine scores: (a) reference genuine scores $\phi(T_1^p, T_2^p)$, where the similarity between two samples of

the same target identity is computed (spoofing index 1, SI_1), and (b) reference genuine scores $\phi(A_1^p, A_2^p)$, where the similarity between two samples of the after-makeup images of a subject is computed (spoofing index 2, SI_2). The two spoofing indices are described below.

Spoofing Index 1: SI_1 is defined as follows:

$$SI_1 = 1 - \min_{i,j} |\phi(A_i^p, T_j^p) - \phi(T_1^p, T_2^p)|, \quad (1)$$

where $i, j \in \{1, 2\}$. Here, we examine if the similarity score between the after-makeup image and the target image is within the range of the score between two samples of the *target identity*. Specifically, $\phi(A_i^p, T_j^p) \approx \phi(T_1^p, T_2^p)$ suggests that spoofing is successful and the output of $SI_1 \approx 1$. For the case $\phi(A_i^p, T_j^p) \ll \phi(T_1^p, T_2^p)$, spoofing is not successful and the output of $SI_1 \approx 0$.

Spoofing Index 2: SI_2 is based on the difference in similarity score between the after-makeup and target images with respect to the after-makeup images:

$$SI_2 = 1 - \min_{i,j} |\phi(A_i^p, T_j^p) - \phi(A_1^p, A_2^p)|. \quad (2)$$

This index is very similar to SI_1 . However, here, we examine if the similarity score between the after-makeup image and the target image is within the range of the score between two *after-makeup samples* of the same person. Specifically, $\phi(A_i^p, T_j^p) \approx \phi(A_1^p, A_2^p)$ suggests that spoofing is successful and the output of $SI_2 \approx 1$. For the case $\phi(A_i^p, T_j^p) \ll \phi(A_1^p, A_2^p)$, spoofing is not successful and the output of $SI_2 \approx 0$.

4. Automated Face Recognition

The following two face matchers were used in this study: a Commercial Off-The-Shelf (COTS) face software and the VGG face matcher.

The VGG face descriptor [20] is computed using a Convolutional Neural Network (CNN) implementation based on the VGG-Very-Deep-16 architecture as characterized in [20]. It is developed based on a triplet-loss training scheme to learn a face embedding that has a similar principle as metric learning.

The effectiveness of these two face matchers was verified using the BLUFR protocol [15] on the LFW dataset³ [10] (as seen in Figure 5). Compared to the original LFW protocol, the BLUFR protocol contains both verification and open-set identification scenarios to fully exploit the potential of the LFW dataset, with focus on low FARs. The results are reported using the “Verification $\mu - \sigma$ Rate” as

³COTS template extraction failed for 113 images of the 12,896 images in the BLUFR protocol. For a fair comparison with VGG, we withheld these images from the VGG matcher as well. Note that the BLUFR protocol only uses 12,896 images of the 13,233 images in LFW.

Table 1. Performance of the COTS and VGG face matchers when using the BLUFR protocol on the LFW dataset. The mean of the Verification Rate over the 10 folds is given by μ and the standard deviation is given by σ .

Matcher	FAR (%)	Verification Rate (%)		
		μ	σ	$\mu - \sigma$
COTS	0.1	81.0	0.647	80.4
VGG Face	0.1	56.6	1.53	55.1

Figure 5. ROC curves averaged over 10 folds of the BLUFR protocol. The “Verification $\mu - \sigma$ Rate” reports the average performance across the 10 folds and subtracts the standard deviation of that performance to yield a lower bound on performance.

described in the BLUFR protocol. Table 1 shows the Verification Rate at 0.1% FAR. Both matchers exceed the baseline performance,⁴ with the COTS matcher surpassing the baseline by a factor of nearly two. Thus, both matchers are good candidates for use in this work.

5. Experiments

In this section, we conduct experiments on the MIFS dataset. By performing these experiments, we seek to understand a) *whether* makeup can be used to spoof faces from the perspective of an automated face matcher, and b) the *extent* to which makeup has the ability to spoof a face recognition system. We reiterate that the subjects represented in the MIFS dataset were not deliberately trying to spoof an automated face recognition system.

5.1. Match Score Analysis

To study the possibility of spoofing, we first analyze the score changes after the application of makeup. For spoofing to be successful, the similarity score between the after-makeup and target images should increase compared to that of the before-makeup and target images. Match scores were generated according to the following protocol:

1. Match B_i^p against T_j^p ($B-T$): the image before makeup is matched against the target image.

⁴<http://www.cbsr.ia.ac.cn/users/sclicao/projects/blufr>

Figure 6. Boxplots displaying similarity score distributions of $B-T$, $A-T$ and $B-A$, as computed by the two face matchers. For each boxplot, the central red line is the median value and the edges of the box correspond to the 25th and 75th percentiles.

Figure 7. Observing the score change between $B-T$ and $A-T$ ($\phi(A, T) - \phi(B, T)$) via histograms for the two face matchers.

2. Match A_i^p against T_j^p ($A-T$): the image after makeup is matched against the target image.
3. Match B_i^p against A_j^p ($B-A$): the image before makeup is matched against the image after makeup.

Figure 6 illustrates the boxplots of score distributions corresponding to $B-T$, $A-T$ and $B-A$. We note an evident upward shift in $A-T$ scores when makeup is applied (compared to $B-T$). This means the application of makeup has increased the chances of a subject being matched against the target image, as reflected by the change in similarity score. Figure 7 visualizes the differences between $A-T$ and $B-T$ as a histogram. As can be seen, the majority of the differences are positive.

5.2. Spoofing Indices

Next, we observe the values of the proposed spoofing indices. As shown in Figure 8, both spoofing indices demonstrate the extent to which makeup has the ability to spoof a particular face recognition system. The higher the value, the more likely that spoofing has occurred. **Both face matchers are observed to be vulnerable to spoof attacks via makeup.**

5.3. Histogram Shift

In addition, we visualize the histograms of match scores that would allow us to see the improvement in makeup induced genuine scores in the context of general genuine and impostor score distributions of the matcher. More specifically, we plot the score distributions of $B-T$, $A-T$, and the LFW dataset in Figure 9. The LFW scores are derived from 13,120 images⁵ in the LFW dataset [10].

In Figure 9(a), we see that the $B-T$ scores fall in the $[0.03, 0.61]$ range and the $A-T$ scores fall in the $[0.13, 0.77]$ range. This indicates that the match scores generally increase after the application of makeup and the histogram shifts to the right. The same phenomenon can be seen in Figure 9(b), but to a lesser extent.

5.4. Identification Experiment

In this setting, the after-makeup and before-makeup images are used as probes, and the target images to be spoofed are placed in the gallery. The goal is to determine whether

⁵113 of the 13,233 LFW images failed the template extraction stage of the COTS matcher. In order for a fair comparison to the OxfordVGG matcher, we withheld these 113 images from the Oxford VGG matcher as well.

Figure 8. Boxplots showing both spoofing indices for the (a) COTS and (b) VGG face matchers. The higher the value, the more likely that spoofing has occurred.

Figure 9. Normalized histogram of similarity scores from the $B-T$ subset, $A-T$ subset, and the LFW dataset for the (a) COTS and (b) VGG face matchers. To improve the visibility of the curves, we crop (b) to the $[0, 0.7]$ range on the x-axis.

target images match at a better rank with after-makeup images than with before-makeup images. We also populate the gallery with 13,120 “background” images from the LFW dataset [10]. We match the probe images against the gallery images and compute the ranks at which the target spoofs are successfully matched. In this context, rank- k denotes the probability that the target identity to be spoofed occurs among the top k matches for a given probe image. The results are summarized in Figure 10, from which we observe that the identification accuracy associated with after-makeup images is significantly higher than that of the before-makeup images. Figure 11 shows examples of subjects being matched at better and worse ranks after the application of makeup. This clearly illustrates that makeup is a viable spoofing method in the context of automated face recognition.

6. Observations and Future Work

In this section we summarize the main findings of this research.

- Based on our experiments with two different face

matchers, we observe that the similarity scores between after-makeup and target face images are higher than those between before-makeup and target-images. This indicates that makeup can be strategically used to spoof a face.

- The identification experiment shows that the target identity being spoofed is likely to match at a better rank after the application of makeup by the impostor.
- Spoofing was successful in some cases and not in others. This is expected as the success of spoofing depends on the source face and target face as well as the makeup procedure used.

Even though subjects in the MIFS dataset were not deliberately launching a spoof attack against an *automated* face matcher, the results presented here demonstrate the potential of using makeup to spoof a face recognition system (see Figure 11). These observations point out the necessity for developing face recognition methods that are less impacted by the application of makeup. There are several ways to potentially address this issue:

(a) COTS Matcher

(b) VGG Face Matcher

Figure 10. Comparison of target identification results before and after the application of makeup. The gallery set consists of the target images to be spoofed as well as 13,120 images from the LFW dataset. The probe set consists of the before-makeup and after-makeup images. The application of makeup for spoofing purposes significantly increases the chance of target identities being matched at higher (*i.e.*, better) ranks.

(a) Target identity retrieved at Rank 26 (before-makeup) and Rank 6 (after-makeup).

(b) Target identity retrieved at Rank 43 (before-makeup) and Rank 1 (after-makeup).

(c) Target identity retrieved at Rank 734 (before-makeup) and Rank 1 (after-makeup).

(d) Target identity retrieved at Rank 1880 (before-makeup) and Rank 1 (after-makeup).

(e) Target identity retrieved at Rank 745 (before-makeup) and Rank 1331 (after-makeup).

(f) Target identity retrieved at Rank 1 (before-makeup) and Rank 13 (after-makeup).

Figure 11. Examples of spoofing attempts with respect to the COTS face matcher. The target images are successfully retrieved at better ranks when using after-makeup images as probes in (a) – (d) (green border) and at worse ranks in (e) and (f) (red border). Note that (c) and (d) contain the same subject but different target identities.

- As indicated in the work of Short et al. [22], polarimetric thermal images are minimally impacted by the application of cosmetic paints. Specifically, their work demonstrated the efficacy of recognizing faces in the presence of makeup by using polarimetric thermal imaging.
- Extracting features that are invariant to the application of makeup can help mitigate the presented type of attack [5]. In addition, makeup detection schemes [3] can be employed for detecting and preprocessing face images prior to face matching.

7. Summary

In this work, we presented preliminary results on makeup-induced face spoofing. We observed in Section 5.1 that similarity scores between a face image and the target face to be spoofed did indeed increase after the application of makeup. In Sections 5.2 and 5.3, we noted that the increase in match score is significant. Finally, in Section 5.4, we showed that a face image is matched with the target identity at a better rank after applying makeup to the former. With the increasing use of face recognition systems in authentication applications, this research suggests that the issue of makeup has to be accounted for in the context of spoof attacks. It is not difficult to envision scenarios where a malicious individual may strategically employ makeup to deceive the system.

References

- [1] A. Anjos and S. Marcel. Counter-measures to photo attacks in face recognition: A public database and a baseline. In *Proc. of IJCB*, 2011.
- [2] M. M. Chakka, A. Anjos, S. Marcel, R. Tronci, D. Muntoni, G. Fadda, M. Pili, N. Sirena, G. Murgia, M. Ristori, F. Roli, J. Yan, D. Yi, Z. Lei, Z. Zhang, S. Z. Li, W. R. Schwartz, A. Rocha, H. Pedrini, J. Lorenzo-Navarro, M. Castrillon-Santana, J. Maatta, A. Hadid, and M. Pietikainen. Competition on counter measures to 2-D facial spoofing attacks. In *Proc. of IJCB*, 2011.
- [3] C. Chen, A. Dantcheva, and A. Ross. Automatic facial makeup detection with application in face recognition. In *Proc. of ICB*, 2013.
- [4] C. Chen, A. Dantcheva, and A. Ross. Impact of facial cosmetics on automatic gender and age estimation algorithms. In *Proc. of VISAPP*, 2014.
- [5] C. Chen, A. Dantcheva, and A. Ross. An ensemble of patch-based subspaces for makeup-robust face recognition. *Information Fusion*, 32(PB):80–92, Nov. 2016.
- [6] A. Dantcheva, C. Chen, and A. Ross. Can facial cosmetics affect the matching accuracy of face recognition systems? In *Proc. of BTAS*, 2012.
- [7] A. Dantcheva and J. Dugelay. Female facial aesthetics based on soft biometrics and photo-quality. In *Proc. of ICME*, 2011.
- [8] M.-L. Eckert, N. Kose, and J.-L. Dugelay. Facial cosmetics database and impact analysis on automatic face recognition. In *Proc. of MMSP*, 2013.
- [9] G. Guo, L. Wen, and S. Yan. Face authentication with makeup changes. *IEEE Transactions on Circuits and Systems for Video Technology*, (99):1–1, 2013.
- [10] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007.
- [11] A. Jain, A. Ross, and K. Nandakumar. *Introduction to Biometrics*. Springer US, 2011.
- [12] N. Kose, L. Apvrille, and J.-L. Dugelay. Facial makeup detection technique based on texture and shape analysis. In *Proc. of FG*, 2015.
- [13] N. Kose and J.-L. Dugelay. Countermeasure for the protection of face recognition systems against mask attacks. In *Proc. of FG*, 2013.
- [14] J. Li, Y. Wang, T. Tan, and A. K. Jain. Live face detection based on the analysis of fourier spectra. In *Proc. of SPIE 5404*, 2004.
- [15] S. Liao, Z. Lei, D. Yi, and S. Z. Li. A benchmark study of large-scale unconstrained face recognition. In *Proc. of IJCB*, Sept 2014.
- [16] J. Maatta, A. Hadid, and M. Pietikainen. Face spoofing detection from single images using texture and local shape analysis. *IET Biometrics*, 1(1):3–10, 2012.
- [17] S. Marcel, M. S. Nixon, and S. Z. Li, editors. *Handbook of Biometric Anti-Spoofing - Trusted Biometrics under Spoofing Attacks*. Advances in Computer Vision and Pattern Recognition. Springer, 2014.
- [18] D. Mee and B. Wong. Medical makeup for concealing facial scars. *Facial Plastic Surgery*, 28(5):536–40, 2012.
- [19] G. Pan, L. Sun, Z. Wu, and S. Lao. Eyeblink-based anti-spoofing in face recognition from a generic webcam. In *Proc. of ICCV*, 2007.
- [20] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *Proc. of BMVC*, 2015.
- [21] K. Patel, H. Han, and A. K. Jain. Secure face unlock: Spoof detection on smartphones. *IEEE Transactions on Information Forensics and Security*, 11(10):2268–2283, Oct 2016.
- [22] N. J. Short, A. J. Yuffa, G. Videen, and S. Hu. Effects of surface materials on polarimetric-thermal measurements: applications to face recognition. *Applied Optics*, 55(19):5226–5233, Jul 2016.
- [23] X. Tan, Y. Li, J. Liu, and L. Jiang. Face liveness detection from a single image with sparse low rank bilinear discriminative model. In *Proc. of ECCV*, 2010.
- [24] Z. Wei, X. Qiu, Z. Sun, and T. Tan. Counterfeit iris detection based on texture analysis. In *Proc. of ICPR*, 2008.
- [25] S. Yoon, J. Feng, and A. K. Jain. Altered fingerprints: Analysis and detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(3):451–464, 2012.