Understanding OSN-Based Facial Disclosure Against Face Authentication Systems

Yan Li, Ke Xu, Qiang Yan, Yingjiu Li, Robert H. Deng
School of Information Systems, Singapore Management University
{yan.li.2009, kexu.2013, qiang.yan.2008, yjli, robertdeng}@smu.edu.sg

ABSTRACT

Face authentication is one of promising biometrics-based user authentication mechanisms that have been widely available in this era of mobile computing. With built-in camera capability on smart phones, tablets, and laptops, face authentication provides an attractive alternative of legacy passwords for its memory-less authentication process. Although it has inherent vulnerability against spoofing attacks, it is generally considered sufficiently secure as an authentication factor for common access protection. However, this belief becomes questionable since image sharing has been popular in online social networks (OSNs). A huge number of personal images are shared every day and accessible to potential adversaries. This OSN-based facial disclosure (OSNFD) creates a significant threat against face authentication.

In this paper, we make the first attempt to quantitatively measure the threat of OSNFD. We examine real-world face-authentication systems designed for both smartphones, tablets, and laptops. Interestingly, our results find that the percentage of vulnerable images that can used for spoofing attacks is moderate, but the percentage of vulnerable users that are subject to spoofing attacks is high. The difference between systems designed for smartphones/tablets and laptops is also significant. In our user study, the average percentage of vulnerable users is 64% for laptop-based systems, and 93% for smartphone/tablet-based systems. This evidence suggests that face authentication may not be suitable to use as an authentication factor, as its confidentiality has been significantly compromised due to OSNFD. In order to understand more detailed characteristics of OSNFD, we further develop a risk estimation tool based on logistic regression to extract key attributes affecting the success rate of spoofing attacks. The OSN users can use this tool to calculate risk scores for their shared images so as to increase their awareness of OSNFD.

Categories and Subject Descriptors

K.6.5 [**Management of Computing and Information Systems**]: Security and Protection—*Authentication*; I.4.9 [**Image Processing and Computer Vision**]: Applications

ASIA CCS'14, June 4–6, 2014, Kyoto, Japan.
Copyright 2014 ACM 978-1-4503-2800-5/14/06 ...$15.00.
http://dx.doi.org/10.1145/2590296.2590315.

Keywords

Face authentication; online social networks; OSN-based facial disclosure

1. INTRODUCTION

Online social networks (OSNs) have been an essential part of modern social life. As the platforms for experience sharing and social interaction, numerous personal data including personal images are being published in OSNs such as Facebook, Google+, and Instagram at every moment. According to a recent report by Facebook, 350 million personal images are published by users on Facebook every day [40]. It is very likely that these images contain facial images where the users' faces can be clearly seen. The large base number indicates that these shared personal images could become an abundant resource for potential attackers to exploit, which introduces the threat of OSN-based facial disclosure (OSNFD).

OSNFD may have a significant impact on the current face authentication systems, which is one of promising biometrics-based user authentication mechanisms. Face authentication have been widely available on all kinds of consumer-level computing devices such as smartphones, tablets, and laptops with built-in camera capability. Popular face authentication systems include Face Unlock [10], Facelock Pro [8], and Visidon [38] on smartphones/tablets, Veriface [24], Luxand Blink [25], and FastAccess [39] on laptops. These systems provide attractive alternatives of legacy passwords, as face authentication requires zero memory efforts from users and usually has higher entropy than legacy password as users tend to choose easy-to-guess passwords [28]. Previously, the major obstacle for an adversary to compromise face authentication is that physical proximity is required to capture a victim's facial images. However, this is no longer necessary since the appearance of OSNFD. OSNFD provides abundant exploitable resources affecting the applicability of face authentication as it compromises its confidentiality, which is one of fundamental requirements for authentication [14, 18]. The facial images used for face authentication are no longer secrets and can be disclosed in large scale due to OSNFD.

In this paper, we make the first attempt to provide a quantitative measurement on the threat of OSNFD against face authentication. We investigate real-world face-authentication systems designed for both smartphones, tablets, and laptops. These systems recognize users by analyzing facial images captured by built-in cameras. Our study collects users' facial images published in OSNs and uses them to simulate the spoofing attacks against these systems. Since all target systems including Google's Face Unlock [10, 8, 38, 24, 25, 39] are closed-source and do not provide any programmable testing interfaces, enormous efforts are made for image collection and testing. We also build a dataset containing important image attributes that are common in real-life photos but rarely used in prior controlled study on face authentication [6, 13].

Our study reveals interesting results indicating that face authentication may not be suitable to use as an authentication factor. Although the percentage of vulnerable images that can be used for spoofing attacks is moderate, the percentage of vulnerable users that are subject to spoofing attacks is high. On average, the percentage of vulnerable users is 64% for laptop-based systems, and 93% for smartphone/tablet-based systems. Our results also show the difference between systems designed for smartphones/tablets and laptops, as smartphones/tablets have to be accessible in more varied environments. Further investigation shows the quality of images is a more important factor affecting the success rate of spoofing attacks compared to quantity. A user who uploads a few clear facial images is more vulnerable than another user who uploads much more facial images of lower quality due to makeup, illumination, or other negative effects. All these findings show that OSNFD has significantly compromised the confidentiality of face authentication.

In order to understand more detailed characteristics of OSNFD, we further develop a risk estimation tool based on our dataset. Logistic regression is used to extract key attributes affecting the success rate of spoofing attacks. It achieves a precision of 81%, a recall of 83%, and an F1 score of 82% on average. It can help users evaluate the risk of uploading an image by calculating a risk score based on the extracted attributes, which makes them aware of the threat of OSNFD.

The contributions of this paper are summarized as follows:

- We investigate the threat of OSN-based face disclosure (OSNFD) against face authentication. Our results suggest that face authentication may not be suitable to use as an authentication factor, as its confidentiality has been significantly compromised by OSNFD.

- We make the first attempt to quantitatively measure the threat of OSNFD by testing real-world face authentication systems designed for smartphones, tablets, and laptops. We also build a dataset containing important image attributes that significantly affect the success rate of spoofing attacks. These attributes are common in real-life photos but rarely used in prior controlled study on face authentication [6, 13].

- We use logistic regression to extract key attributes that affect the success rate of spoofing attacks. These attributes are further used to develop a risk estimation tool to help users measure the risk score of uploading images to OSNs.

2. PRELIMINARIES

2.1 Face Authentication

Face authentication is a biometrics-based user authentication mechanism, which verifies a user's identity by using information extracted from the user's facial features. As illustrated in Figure 1, a typical face authentication system uses a camera to capture the user's facial image/video as input, and then verifies it with enrolled biometric information for the claimed identity. The objective of a face authentication system is to recognize a user as long as the input is collected from the legitimate user, while rejecting the inputs from all other users.

Two key modules are involved in this verification process. The first module is the face detection module, which identifies the face region and removes irrelevant information of an image. The processed image is then passed to the next module named face matching. This module computes a similarity score for the input image based on an enrolled face template containing key features which can be used to distinguish a user from other users and imposters. Different algorithms may be used for these two modules, but all face authentication systems generally have these two modules and follow this work flow. In the end, a face authentication system outputs the final decision (i.e. accepting or rejecting a claim) according to whether or not the similarity score is higher than a matching threshold. This threshold is carefully chosen so as to achieve a proper balance between false rejection rate and false acceptance rate.

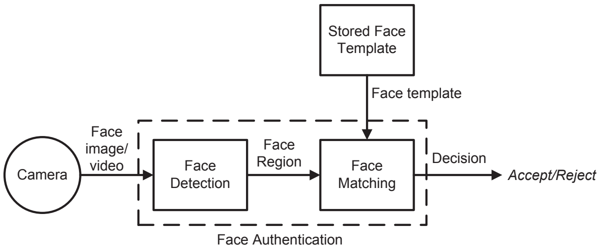

Figure 1: Work flow of a typical face authentication system

2.2 OSN-based Facial Disclosure and Threat Model

The OSN-based facial disclosure (OSNFD) addresses the issue when users' face biometrics are involuntarily disclosed by sharing personal images in OSNs. These disclosed face biometrics would raise security risks against face authentication systems.

It is a well-known limitation of face authentication that it is subject to spoofing attacks based on captured face biometrics, where an adversary attempts to circumvent user authentication by replaying a victim's facial images/videos collected at an early time. As shown in Figure 1, a face authentication system is not expected to tell whether an input image is from a live user or from a captured image/video, as they are all valid inputs from a legitimate user collected at different times. Nevertheless, the impact of these attacks was believed to be limited due to the requirement that an adversary had to be physically close to a victim in order to collect the required information. Therefore, it is generally considered sufficiently secure as an authentication factor for common access protection [5], as we observe that many face authentication systems [10, 8, 38, 24, 25, 39] such as Google's Face Unlock and Lenovo's Veriface, are widely available on all kinds of consumer-level computing devices. Considering its zero-memory requirement, it does provide an attractive alternative for legacy passwords.

However, this belief may be questionable since OSNFD becomes a common phenomenon. OSNFD supplies an adversary with abundant facial images to exploit and makes large-scale identity theft possible for those who use face authentication. Our work investigates the OSNFD threat and quantitatively measures its impacts. We consider OSNFD-based attacks where an adversary attempts to forge a valid input from image resources disclosed from OSNFD so as to pass face authentication. Our study focuses on image-based attacks unless explicitly mentioned.

The OSNFD threat may be mitigated with liveness detection technologies, which rely on extra information sources or heuristic algorithms to distinguish a live user from a captured image/video. All the existing sophisticated liveness detection technologies associate with considerable costs, which will be explained later in Section 5.2. This may explain that only weak liveness detection technologies are currently deployed on the face authentication systems designed for consumer-level computing devices [29, 19]. For example, eye blinking detection is a common heuristic used by many face authentication systems [10, 38, 29] including Google's Face Unlock; however, it can be easily bypassed using two facial images as demonstrated in [31]. Similar tricks can also apply to other weak liveness detection mechanisms such as head rotation detection [29, 31]. Even worse is that the existing liveness detection mechanisms are disabled by default in most popular face authentication systems [10, 8, 38, 24, 25, 39], as they may have negative impacts on accessibility.

3. DATA COLLECTION AND EMPIRICAL ANALYSIS

In order to quantitatively measure the impacts of OSNFD, we conduct a user study to collect real personal images that have been shared in OSNs. The collected images are used to test against real-world face-authentication systems chosen from the most popular face authentication products in terms of user base [37, 11]. This section describes the detailed process of data collection and the results of our empirical analysis. We use the following classifications in our discussion.

First, we classify the security settings of a face authentication system into *low* and *high*. Most of face authentication products [10, 8, 38, 24, 25, 39] provide very limited choices on security settings that generally affect the recognition threshold used in the face matching module. For example, Google's Face Unlock [10] does not provide any option for users to adjust its security strength. Most of our tested products [8, 38, 24, 25, 39] only have two options for users, labeled as "high accessibility" (i.e. low security) and "high security". Only Lenovo's Veriface [24] provides a scroll-bar for users to adjust its security strength from the lowest to the highest. Therefore, we use "low" to indicate that a target system enforces the weakest security protection, and use "high" to indicate the strongest security protection achievable to the system.

Second, we classify face authentication systems into *mobile* and *traditional*. A system is labeled as mobile if it is used for smartphones or tablets, while a traditional system is used for laptops or desktops. A mobile system is usually more tolerant to varied environments, as it should be accessible no matter where a user uses the device. Laptops is considered as traditional as it is not expected to be used from anywhere at any time like what users expect smartphones and tablets.

Third, we classify users into different groups according to the pattern of their sharing behaviors. As observed in our study, it is quite common that a user tends to upload edited images where facial landmarks are significant changed to create better visual appeal. Therefore, it is also an important factor that needs to be considered.

These classifications represent three major factors that affect the effectiveness of OSNFD-based attacks, which are security settings, target platforms, and user behaviors, respectively. We use them as controlled parameters to evaluate the severity of OSNFD, and more sophisticated statistical analysis will be given in the next section to identify the key attributes that can be used to mitigate the OSNFD threat.

3.1 Data Collection

There are 74 participants involved in our study including 36 males and 38 females with age range between 19 and 35. Most of these participants are students in our university. Each participant is paid with 10 dollars as compensation. The study is conducted in a quiet room. The study consists of three parts. In the first part, we ask each participant to select and download 20 *facial* images published within the last 12 months in popular OSNs such as Facebook, Google+, Instagram, etc. A facial image is defined as an image where a participant's face can be seen. But the participant's face may be affected by many negative effects such as blur, occlusion (e.g. covered by a sunglasses), head rotation (e.g. non-frontal head pose). All these effects will be examined in our study.

In the second part, we capture the participant's facial images with 35 controlled head poses and 5 facial expressions using a Canon EOS 60D (18.0-megapixel DSLR CMOS camera). The resulting images are 5184×3456 in size with inner pupil distance of the subjects typically exceeding 400 pixels. 35 controlled head poses are specified by both horizontal and vertical rotation. Rotation angles are represented as (rot_H, rot_V) where rot_H corresponds to the angle of horizontal rotation while rot_V corresponds to the angle of vertical rotation. The value range of rot_H contains 0°, 10° to left/right, 20° to left/right, 30° to left/right while the value range of rot_V contains 0°, 10° to up/down, 20° to up/down. We choose these boundary values according to the common restriction of existing face authentication systems [1], where a participant should not pass user authentication if rot_H exceeds 30° or rot_V exceeds 20° degrees. On the other hand, 5 facial expressions include neutral expression, smile without showing teeth, smile showing teeth, closed eyes, and open mouth. Continuous lighting system is used to eliminate the shadow on the participants' faces.

We use a helmet equipped wit a gyroscope to control head rotation of the participants. The use of gyroscope has advantages over the other approaches, which includes attaining theoretical accuracy of less than 1 degree, ignoring the head position, measuring only orientation, not affected by metallic interference [27]. For each head pose, we firstly ask the participants to face to the DSLR camera and help them adjust their heads to frontal position in the way similar to [13]. Then the participants rotate their heads to the required angles with help of the gyroscope. The gyroscope generates real-time rotation angles and broadcasts them via WiFi. This rotation information will be received and displayed on an iPad screen, and shown to the participants. Thirdly, we ask the participants to hold their head poses and one of our researchers then removes the helmet gently and quickly in order to avoid movement of the heads during helmet removal. After that, the images of each head pose are captured immediately.

In the final part, the participant will be asked to fill in a questionnaire for collecting the participant's attitudes towards usage of face authentication systems and sharing behaviors in OSNs.

3.2 Empirical Results

Based on the collected images, we inspect the realistic threat of OSNFD against the latest version of popular real-world face authentication systems. We use the common experiment procedure similar to prior work [22, 6], which is described as follows: The frontal image is first used to enroll each participant into a face authentication system. Then we use a participant's own OSN images to test whether it can be used to log in a target face authentication system for his/her own account. The participant's OSN images are displayed on an LCD screen with resolution 1600×900 pixels, and the result whether a target system can be spoofed by an OSN image will be recorded for each system and each image.

Our analysis uses two basic metrics, namely *vulnerable images* and *vulnerable users*. A vulnerable image, denoted by $VulImage$, is defined as a facial image which is wrongly accepted as a genuine user by a face authentication system during user authentication and therefore enables an adversary to circumvent the face authentication system. A vulnerable user, denoted by $VulUser$, is a user enrolled in a face authentication system who has at least one vulnerable image published in OSNs.

Table 1 shows that the face authentication systems are vulnerable to the OSNFD in general. On average, 39% of the OSN images and 77% of the participants are vulnerable. Among popular face authentication systems, Visidon is more vulnerable in low security level, for which 68% of the images and 97% of the participants are vulnerable. Especially for Google's Face Unlock that comes as a built-in feature of all Android-based systems whose version is higher than 4.0 [10], 45% of the OSN images and 86% of the participants are vulnerable.

Table 1: Overall percentage of $VulImage$ and $VulUser$

	$VulImage\%$	$VulUser\%$
Face Unlock	45%	86%
Facelock Pro	46%	96%
Visidon	68%	97%
Veriface	27%	73%
Luxand Blink	20%	41%
FastAccess	33%	80%
Average	39%	77%

Although the percentage of vulnerable images is moderate, the quantity of the vulnerable images is large due to the huge amount of images in OSNs. These large amount of vulnerable images create resources online for potential attacks. Even worse, users share their personal images with their friends in OSNs, most of them tend to publish the images where the users' faces can be clearly viewed for easier recognition. Consequently, the percentage of vulnerable users would be high as observed in our study. The following subsections will further analyze the detailed characteristics of these vulnerable images and users from three major perspectives, security settings, target platforms, and user behaviors.

3.2.1 Impacts of Security Settings

Security settings specify the security strength of a face authentication system against potential attacks. As previously explained, most of face authentication products [10, 8, 34, 23, 24, 35] provide very limited choices on security level. So we focus our analysis on lowest and highest security level that can be provided by each system, which are denoted as low security and high security, respectively. Since there is only one security level in Face Unlock and the observed security strength of Face Unlock is comparable to the other systems in low security level, we classify its security level as low. As expected, Figure 2 shows that the face authentication systems in low security level are facing more severe OSNFD threat than those in high security level. On average, 40% of the images and 79% of the participants are vulnerable for the face authentication systems in low security level while 8% of the images and 30% of the participants are vulnerable for the face authentication systems in high security level.

The change of security settings generally affects the recognition threshold in the face matching module. As the security level is raised, the recognition threshold becomes higher which imposes more restrictions for matching between login facial image and prestored facial image. Therefore the face authentication imposes more rigid restrictions on the login facial image. The major restrictions observed in our study are head pose and lighting condition.

For head pose, we use *acceptable head pose range* to measure the tolerance of a face authentication system on head pose variations. It describes the head rotation range of head poses with which at least 50% of the participants successfully log in the face authentication systems. In these tests, we use participant's frontal image for enrollment and use the images collected with controlled head poses as test inputs (i.e. login images). Figure 3 shows the average results computed from all tested systems, where each closed curve corresponds to the acceptable head pose range. The results for each individual system that indicate the difference between high security and low security are similar to Figure 3, which are not shown.

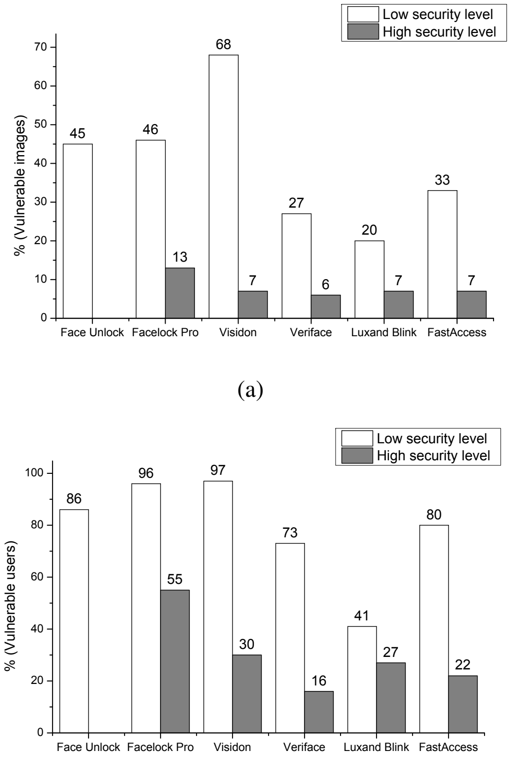

(a)

(b)

Figure 2: Percentage of $VulImage$ and $VulUser$ in different security levels

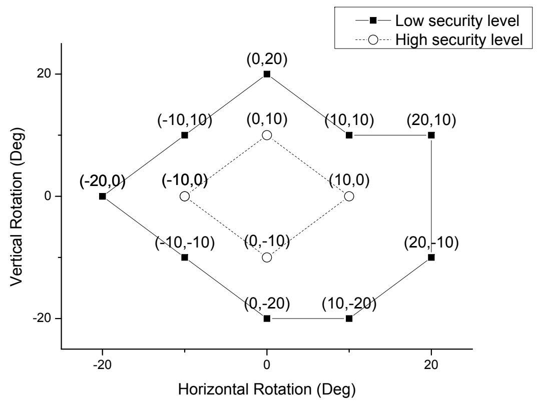

Figure 3: Tolerance of the rotation range of head pose

For lighting condition, we further classify it into different types of illumination and low lighting [9, 43, 20]. The face authentication systems in low security level are observed to have higher tolerance for variation of lighting conditions than the systems in high security level. In our study, illumination is observed in 27% (394 out of 1440) of the OSN images while low lighting is observed in 18% (266 out of 1440) of the OSN images. On average, 81% of the OSN images with illumination and 79% of the OSN images with low lighting cannot be used to log in the face authentication systems in low security level while 96% of the OSN images with illumination and 94% of the OSN images with low lighting cannot be used to log in the systems in high security level.

On the other hand, a face authentication system in low security level has higher tolerance for varied login environments, which is necessary for the system to be usable in the complex environments. As a tradeoff for higher security strength, the false rejection rates in high security level may be significantly increased. As shown in the follow-up experiment described in Section 5.1, the false rejection rate could be as high as 85%. This will cause a significant concern on the accessibility. From our questionnaire on user perception, 70% of the participants think it is important to successfully log in their smartphones, tablets, or laptops at the time they want to use. If the face authentication system is not always functional, 67% of the participants give up using the system which causes the serious accessibility problem to their devices. This may also explain why the popular face authentication systems always use low security level by default.

3.2.2 Impacts of Target Platforms

The target platform of a face authentication system imposes the platform-specific requirements on both security and usability. In our tested systems, Face Unlock, Facelock Pro, and Visidon are targeting for mobile platform, while Veriface, Luxand Blink, and FastAccess are targeting for traditional platform.

Figure 4 shows that the OSNFD threat for mobile platform is generally more severe than the OSNFD threat for traditional platform. On average, in low security level, 53% of the images and 93% of the participants are vulnerable for the face authentication systems on mobile platform while 27% of the images and 64% of the participants are vulnerable for the systems on traditional platform. In high security level, 10% of the images and 43% of the participants are vulnerable for the face authentication systems on mobile platform while 7% of the images and 22% of the participants are vulnerable for the face authentication systems on traditional platform.

These results clearly show the difference caused by platform-specific requirements. Compared to a traditional system, a mobile system is usually designed to be more robust and more tolerant to varied environments such as outdoor environment in order to meet accessibility expectation by users. Meanwhile it leads to the more severe OSNFD threat for mobile platform based systems. This difference is confirmed by the results of our questionnaire, which shows that 91% of the participants believe that it is important to log in smartphones or tablets in both indoor and outdoor environment while only 36% of the participants think it is important to log in laptops in both indoor and outdoor environment.

This difference is also revealed in our tests on head pose and lighting condition. Figure 5 shows the face authentication systems targeting for mobile platform have higher tolerance for variations of the head poses than the systems targeting for traditional platform.

Our tests on lighting conditions further show the face authentication systems targeting for mobile platform are more tolerant to variations of the lighting conditions. In our study, 81% of the OSN images with illumination and 77% of the OSN images with low lighting cannot be used to log in the face authentication systems targeting for mobile platform, while these rates increase to 96% for the images with illumination and 96% for the images with low lighting on traditional platform.

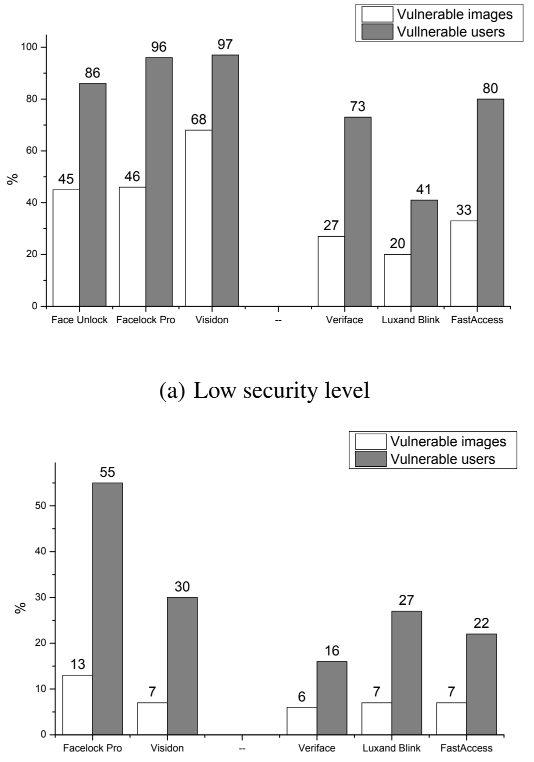

(a) Low security level

(b) High security level

Figure 4: Difference in $VulImage$ and $VulUser$ between systems targeting for mobile platform and traditional platform.

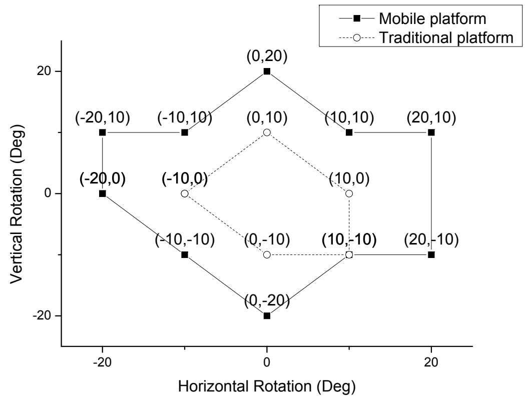

Figure 5: Difference in the tolerance of the rotation range of head pose.

3.2.3 Impacts of User Behaviors

The difference in user behavior is another major factor influencing the quality of shared images that decides whether these images can be eventually used for successful OSNFD-based attacks. Our study reveals that the participants who publish more facial images in OSNs are not necessarily more vulnerable than those who publish less facial images in OSNs. In fact, the OSNFD threat is more severe among the participants who publish facial images with higher quality in OSNs.

To illustrate the impact of user behaviors, we use the different sharing behaviors and the different OSNFD threat between females and males as example. In our study, female participants are reported to publish facial images in OSNs more frequently than male participants in general. On average, each of the female participants publishes 65 facial images per year while each of the male participants publishes 34 facial images per year. However, the OSNFD threat for the females is less severe than that for the males, as shown in Figure 6.

This can be explained by the lower quality of the OSN images published by the females. We find that the female participants are more likely to publish blurred images, edited images, or images with their makeup. The blur, edit, and makeup can degrade the quality of an image and therefore lead to the difficulty in face recognition [16, 7]. In our study, 12% of the OSN images suffer from these negative effects. Among these low quality images, 61% are published by the females while only 39% of the images are published by the males. All of these blurred, makeup, or edited images fail to pass at least one face authentication system.

4. STATISTICAL ANALYSIS AND RISK ESTIMATION

Although the OSNFD threat is significant as shown in the previous section, we observe the effectiveness of OSNFD-based attacks may be significantly reduced by manipulating certain attributes of facial images. In this section, we extract these key attributes via statistical analysis and use them to develop an estimation tool for end users to calculate the risk of their shared images.

4.1 Key Attributes Affecting OSNFD-based Attacks

From the theoretical perspective, there are still many challenges for face recognition algorithms. These challenges also become key attributes that limit the effectiveness of OSNFD-based attacks. The common attributes addressed in the prior study [1] include head pose, lighting condition, facial expression, facial occlusion, and image resolution. Beside these traditional attributes, we also observe blur, facial makeup, and editing (using Photoshop-like software) as the extra key attributes which often appear in the real world images shared in OSNs, though they are usually not considered in the controlled settings of traditional study on face authentication. We describe the details of these key attributes as follows.

Head pose is a prominent challenge to face recognition. The performance of face recognition algorithms in face authentication can be significantly affected if the head pose in a login image and the head pose in the pre-stored facial image are different [43]. The affecting variations of a head pose mainly include two out-of-plane rotations, namely horizontal rotation and vertical rotation [27].

Lighting condition is another prominent challenge in the realm of face recognition. The variation of lighting conditions mainly includes illumination and low lighting [9, 43, 20]. The illumination is mainly caused when direct light shoots on the 3D structure of a face and strong shadows can be casted which diminish facial features [9, 43]. The illumination can be classified into side illumination and top/bottom illumination [9]. Low lighting is another negative lighting condition, which usually happens when a facial image is taken in dim environment or with extreme bright background. The low lighting may diminish facial features since the luminance in face region is too low for face recognition algorithms to recognize [20].

Facial expression such as smile, surprise, etc, can change face geometry and therefore affect the performance of face recognition algorithms [1]. The common facial expressions include neutral expression, smile without showing teeth, smile showing teeth, closed eyes, open mouth, and other expressions.

Facial occlusion often happens in real world due to additional accessories on face, such as sunglasses, scarf, hands on face, etc. The occlusion can result in the failure of face appearance representation or imprecise facial feature searching and localization, and therefore have negative influence on the performance of face recognition algorithms. The common facial occlusions include forehead occlusion, eyebrow occlusion, eye occlusion, cheek occlusion, and mouth occlusion [1].

The **resolution** of an image can affect accuracy of facial landmark localization and therefore influence the performance of face recognition algorithms. As the resolution of face images decreases, the performance of the face recognition algorithms drops [43].

The **blur** in a facial image causes difficulty in accurate localization of edges of facial region and facial landmarks (i.e. eyes, nose, mouth, etc) by face recognition algorithms and therefore harms the performance of the algorithms.

Facial **makeup** can substantially change the appearance of a face and facial landmarks, such as the alternations of perceived facial shape, nose shape, location of eyebrows, etc. These alternations by the facial makeup, especially by non-permanent facial makeup, challenge face recognition significantly [7].

The **editing** of an image introduces noise pixels and change the appearance of the face in the image [2, 7]. Face recognition algorithms can be affected by these noises and appearance changes due to the edited image.

All these attributes significantly degrade the image quality and therefore lead to the failure of OSNFD-based attacks. They are used as input parameters to build our risk estimation tool in the next section.

4.2 Risk Estimation Model

We use binomial logistic regression [15] to model the impact of the key attributes introduced in the previous subsection. The notions of these attributes are defined in Table 2. Then the key attributes of each image can be represented by an input parameter vector, denoted as $V = (rot_H, rot_V, ill_{sd}, ill_{tb}, dm, bg, FEx_n, FEx_s, FEx_{st}, FEx_{ce}, FEx_m, FEx_{other}, Occ_{fh}, Occ_{eb}, Occ_{eye}, Occ_{chk}, Occ_{mh}, res, blur, mk, ed)$.

For the output, we assign an OSN image to either a positive class or a negative class. The positive class means the image can be used to pass the login of a specific face authentication system, otherwise the image will be in the negative class.

Binomial logistic regression is a classic probabilistic classification model [15], which accepts multiple predictor variables as inputs, and predicts the outcome for a dependent variable which has only two possible types, such as “positive” vs “negative”. Thus it is a proper tool to calculate the probability of an image assigned to the positive class based on the key attributes extracted from an OSN image. Given a parameter vector V_i of a facial image i and a face authentication system in a security level, the regression function is

$$\ln\left(p_i/(1-p_i)\right) = \beta_0 + \beta_1 v_1 + \cdots + \beta_m v_m \qquad (1)$$

where p_i is the probability that an image i is assigned to the positive class, v is a parameter in V_i, and β is a regression coefficient. The risk score of the facial image i is the value of p_i. The facial image

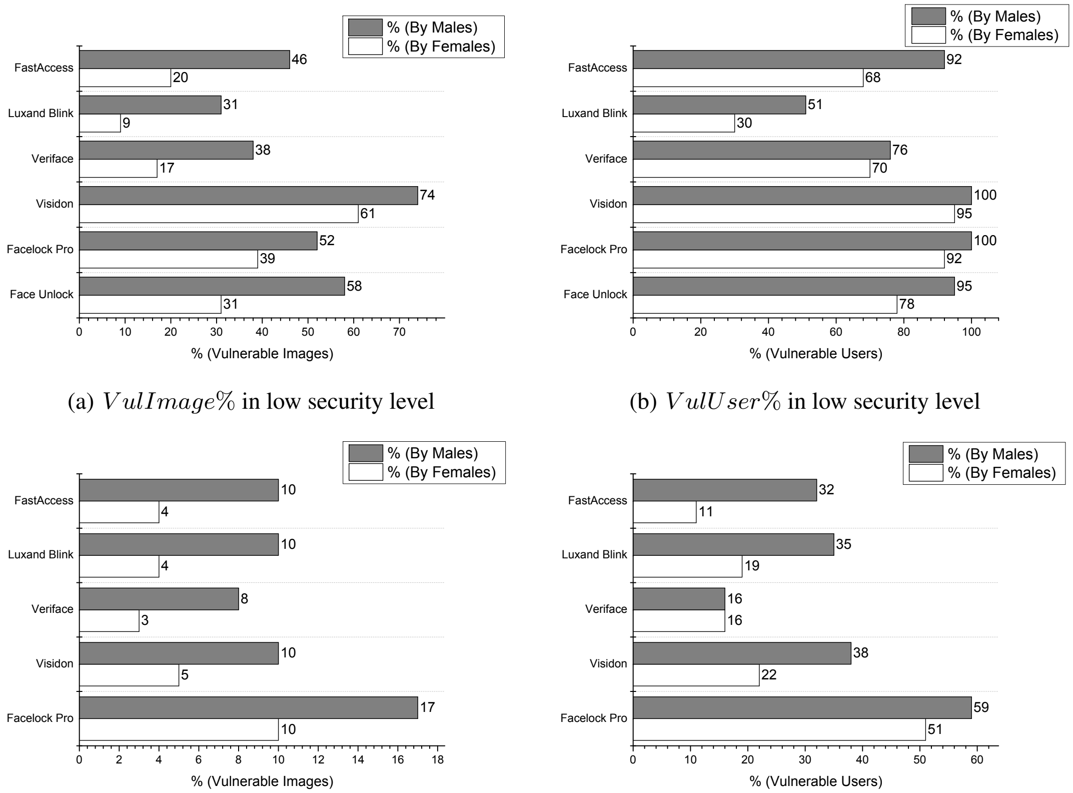

(a) $VulImage\%$ in low security level

(b) $VulUser\%$ in low security level

(c) $VulImage\%$ in high security level

(d) $VulUser\%$ in high security level

Figure 6: Difference in $VulImage$ and $VulUser$ between females and males

i is assigned to the positive class if $p_i \geq 0.5$. Otherwise, i is assigned to the negative class. The correctness of these assignments is verified with the ground truth data collected from the previous empirical analysis.

Table 2: Parameters related to the key attributes

Attribute	Parameter	Notation
Head pose	Horizontal rotation	rot_H
	Vertical rotation	rot_V
Lighting condition	Side illumination	ill_{sd}
	Top/bottom illumination	ill_{tb}
	Dimness	dm
	Bright background	bg
Facial expression	Neutral	FEx_n
	Smile without showing teeth	FEx_s
	Smile showing teeth	FEx_{st}
	Closed eyes	FEx_{ce}
	Open mouth	FEx_m
	Other expressions	FEx_{other}
Facial occlusion	Occluded forehead	Occ_{fh}
	Occluded eyebrow	Occ_{eb}
	Occluded eye	Occ_{eye}
	Occluded cheek	Occ_{chk}
	Occluded mouth	Occ_{mh}
Resolution	Resolution	res
Blur	Blur	$blur$
Facial makeup	Makeup	mk
Edit	Edit	ed

For each combination of face authentication system and its security level, we examine the model fitting of binomial logistic regression and the significance of the parameters by using the real world OSN images and run binomial logistic regression on SAS software [33]. The likelihood ratio test and wald statistic [15] for all the face authentication systems are smaller than 0.0001.

Our statistical analysis shows the most influential attributes are resolution res, occluded eye Occ_{eye}, makeup mk, and illumination ill_{sd}. Resolution res has positive impact on the risk of OSNFD. It is because higher resolution contributes to more accurate facial landmark localization and results in better performance of face recognition and increases the risk of OSNFD. The occluded eye Occ_{eye}, makeup mk, and illumination ill_{sd} have negative impact and lower the risk of OSNFD. In particular, the occluded eye leads to decrease in the performance of face recognition algorithms, as accurate localization of eyes is important for the alignment process in all major face recognition algorithms [1]. Makeup can significantly change the appearance of the face and the facial landmarks and therefore lowers the performance of face recognition. The illumination is a prominent attribute which causes difficulty in face recognition since it diminishes facial features.

The parameters related to other attributes, including head pose and facial expression, are generally not statistically significant. Among the collected OSN images, the variations of head pose and facial expression are limited since users are usually cooperative when these images are captured and tend to publish the images from which they are easily recognized. As observed in our study, the head poses in most OSN images are within the acceptable head pose ranges of the face authentication systems, which causes the insignificance due to lack of samples with extreme head pose. On the other hand, facial expressions observed in most OSN images are only mild-mannered expressions including neutral expression, smile without showing teeth, smile showing teeth, closed eyes, open mouth. These common expressions do not have significant impact as they have been well handled in current face recognition algorithms [1]. Other extreme facial expressions, such as making faces, do significantly affect the face recognition, but they are observed in only 5% of the OSN images.

4.3 Model Evaluation

To evaluate the performance of the proposed risk estimation tool, we use cross-validation method. In each round, for each of the face authentication systems in a specific security level, we randomly choose 80% of the OSN images to train the model and use the risk estimation tool to automatically classify the rest of the images. The above process is repeated by 10 rounds. The performance is measured by standard classification evaluation metrics, including precision, recall, and F1 score [32].

Precision is defined as the percentage of the true positive images among the images assigned to the positive class by the risk estimation tool, which can be calculated by $tp/(tp + fp)$ where tp is the number of true positive images and fp is the number of false positive images. Recall is defined as the percentage of the true positive images detected by the risk estimation tool among the positive images in ground truth, which can be calculated by $tp/(tp + fn)$ where tp is the number of true positive images and fn is the number of false negative images. F1 score considers both the precision and the recall, which can be calculated by $\mathsf{F1} = 2 \times \mathsf{precision} \times \mathsf{recall}/(\mathsf{precision} + \mathsf{recall})$.

Table 3 shows the performance evaluation metrics of the risk estimation tool. On average, the risk estimation tool achieves a precision of 81%, a recall of 83%, and an F1 score of 82%. The performance evaluation indicates that the risk estimation tool detects most of the vulnerable images which can lead to successful OSNFD-based attacks if these images are published in OSNs.

Table 3: Effectiveness of our risk estimation tool

System	Security level	Precision	Recall	F1 score
Face Unlock	N/A	73%	77%	75%
Facelock Pro	Low	70%	69%	69%
	High	81%	75%	78%
Visidon	Low	79%	90%	84%
	High	86%	92%	89%
Veriface	Low	79%	68%	73%
	High	90%	98%	94%
Luxand Blink	Low	84%	87%	85%
	High	87%	90%	88%
FastAccess	Low	77%	67%	72%
	High	89%	95%	92%
Average	N/A	81%	83%	82%

5. DISCUSSION

5.1 Tradeoff between Security and Accessibility

Clear tradeoffs between security and accessibility can be observed in our tested systems, which are decided by security settings and target platforms as analyzed in Section 3.2. The increasing security strength inevitably decreases the accessibility. We conduct a follow-up experiment to collect quantitative evidence for the impact of these tradeoffs.

20 participants from the main user study are invited for this follow-up study. The participants need to enroll their faces in the 6 face authentication systems in low/high security level in a meeting room with normal lighting, respectively. To mimic the different login environment, the experiments are conducted between 2pm-4pm in a sunny day at four fixed indoor/outdoor locations, including 1) a meeting room with normal lighting condition, 2) a meeting room with dim lighting condition, 3) outdoor ground in the sunshine, and 4) shelter of building. This setting simulates a situation when a user registers in one place, but tries to access the system in many other places. The participants are asked to login by using each face authentication systems. In this experiment, there are no OSN images, but only live legitimate users who attempt to access a face authentication system. Each participant has at most three attempts for each login before we record it as a false rejection.

Table 4 shows the false rejection rates of the face authentication systems in low security level are lower than those of the face authentication systems in high security level in overall. Moreover, the face authentication systems on mobile platform have lower false rejection rates than those on traditional platform. The highest observed false rejection rate is 85% for Veriface in high security level. This accessibility degradation could be a disaster for end users. In our questionnaire, 91% of the participants believe that it is important to log in smartphones and tablets in both indoor and outdoor environments, while 36% of the participants think that it is important to log in laptops in both indoor and outdoor environments. If a face authentication system is set to high security level in order to mitigate the OSNFD threat, the system will be less tolerant for complex environments and violate the users' need of accessibility.

Table 4: Significant increase in false rejection rates when using high security level settings. The increments of false rejection rates are more significant for traditional platform-based systems (the last three systems).

System	Security level	Room+ normal lighting	Room+ dim lighting	Outdoor ground	Shelter
Face Unlock	N/A	0%	5%	10%	0%
Facelock Pro	Low	0%	10%	10%	0%
	High	0%	45%	60%	25%
Visidon	Low	0%	5%	5%	0%
	High	5%	55%	65%	50%
Veriface	Low	0%	25%	35%	20%
	High	10%	60%	85%	60%
Luxand Blink	Low	0%	30%	50%	45%
	High	5%	55%	70%	55%
FastAccess	Low	0%	15%	30%	15%
	High	5%	55%	65%	55%

5.2 Costs of Liveness Detection

Liveness detection could be a mitigation for OSNFD-based attacks, which is designed to distinguish between a live face and a facial image in front of the camera. The most common liveness detection mechanisms deployed on popular face authentication systems are eye-blinking and head rotation detection, as they have the advantages of no additional hardware support, requiring moderate image quality, and involving relatively low usability cost. This is important to all consumer-level products that are price-sensitive and accessibility-first. However, these two mechanisms can be easily bypassed with one or two pre-catched images as shown in [31]. The practicality of these attacks is also verified by our experiments.

Besides these two simple mechanisms, several sophisticated liveness detection techniques have been proposed for face authentication. However, all of them are associated with considerable costs as shown in Table 5 [29]. Their costs include requiring additional hardware, high quality images, ideal environment that are usually not universally available, and high user collaborations that may cause inconvenience. This indicates they may not be suitable for consumer-level face authentication systems. It still remains a challenge to deploy reliable and practical liveness detection in face authentication systems that can be used by the public.

Table 5: Costs associated with existing liveness detection mechanisms for face authentication. * sign indicates a requirement involves a significant cost for end users or device manufacturers.

Liveness detection	Image quality	Additional hardware	Usability cost
Eye blinking	Low	No	Low
Mouth movement	Middle	No	Middle
Degradation	High*	No	Low
Head movement	High*	No	Middle
Facial expressions	High*	No	Middle
Facial thermogram	N/A	Yes*	Low
Multi-modal	N/A	Yes*	Middle/High*
Facial vein map	N/A	Yes*	Middle
Interactive response	N/A	Yes*	High*

5.3 Implications of Our Findings

Face authentication does provide an attractive alternative of user authentication for its non-intrusive and zero-memory procedure. However, the appearance of OSNFD brings a significant threat to question the practicality of face authentication as a usable authentication factor. Nowadays, a huge amount of personal facial images/videos have been published in OSNs that can be accessible to potential adversaries without the previously required physical proximity. Therefore, face biometrics can now be disclosed in large scale and acquired by adversaries remotely. Face biometrics are no longer secrets only owned by the users and can be disclosed to anyone who has access to victim's personal images shared in OSNs.

Raising the security level of face authentication systems could mitigate the OSNFD threat by scarifying the accessibility, which leads to the inconvenience for legitimate users. Liveness detection is another major countermeasure to mitigate the spoofing attack against the face authentication systems. Unfortunately, existing liveness detection techniques available on consumer-level computing devices can be easily circumvented by one or two images. More reliable liveness detection like multi-modal mechanisms usually relies on using additional authentication factor (e.g. another biometrics such as voice and fingerprint). This introduces another liveness detection problem for the additional authentication factor, which may not be reliable. For example, voice and fingerprint can also be spoofed. Even worse, more serious privacy concerns will rise if a system requires to collect many biometrics information from a user [42], which may eventually cause the rejection of the liveness detection mechanism.

As the emergence of OSNFD, the face biometrics is losing confidentiality which is one of the fundamental requirements for a usable authentication factor. Moreover, the existing liveness detection techniques are either too weak to defend against the OSNFD or too difficult to be deployed on the consumer-level devices. All these findings suggest that face authentication may not be a proper authentication factor unless we can resolve the discovered problems.

5.4 Limitations

Ecological validity is a challenge to any user study. Like most prior research [12, 35, 3], our study only recruits students in university. These participants are more active in using consumer-level computing devices and sharing images in OSNs. Thus the evaluation of the OSNFD may vary with other populations.

In the user study design, it is still a challenge to collect facial images with precisely controlled head poses [27]. Like the prior head pose data sets [13, 23, 34], the accuracy of the head poses in our data set may be affected by the poor ability of the participant to accurately direct his/her head, the unconscious movement of human beings and limit of resources. In another experiment of examining the false rejection rates of the face authentication systems, we choose 4 locations to mimic different login environments in daily life. Since it is impossible for all the participants to do the tests at the same time and at the same physical positions, the background of image inputs captured by the camera may change.

Another challenge in our study is to accurately estimate parameters [21] such as head pose, illumination, and makeup in our collected OSN dataset. Since the accuracy of automatic labeling tools is limited [1, 30], we manually label the OSN images with the help of automatic tools and follow the similar validating methodology used in prior study [21, 17, 41]. For each OSN image, we estimate the head pose with typical head pose estimation algorithms including POSIT and LGBP [27]. And we manually validate the estimation of the head pose by comparison between the OSN image and the participant's images with controlled head poses. We manually label the parameters related to lighting conditions according to the shadow and histogram of face region similar to the approaches in [21, 17]. The parameters related to facial expressions are label by comparing the OSN image with the images captured in our user study, which is similar to [21, 17]. We use popular face detection software Picasa to mark the face region with a rectangle in the image and calculate the resolution of the face region. The parameters related to the attributes of blur, makeup, and edit are labeled in the way similar to [21, 17, 41].

It is also possible to further improve our risk estimation tool. To our best knowledge, our work is the first attempt to semi-automatically detect the vulnerable images that can be used to attack face authentication. Our current risk estimation tool can serve as a baseline for future improvement by refining the key parameters and the statistical model. It is also valuable to incorporate automatic high accuracy labeling for those hard-to-label attributes like illumination and facial makeup, once the ongoing research [27, 9, 7] resolves these challenges.

6. RELATED WORK

In this section, we summarize the closely related work in terms of face recognition, spoofing attack, and liveness detection.

In face recognition, holistic approaches and local landmark based approaches are the two major types of popular face recognition algorithms [1, 43]. The holistic approaches, such as PCA-based algorithms and LDA-based algorithms, use the whole face region as input. Local landmark based approaches extract local facial landmarks such as eyes, nose, mouth, etc and feed locations and local statistics of these local facial landmarks into a structure classifier.

Face authentication is an important application of face recognition, which validates a claimed identity based on comparison between a facial image and an enrolled facial image and determines either accepting or rejecting the claimed identity [26]. Trewin et al. [36] show that the face authentication is faster and causes lower interruption of user memory recall task than voice, gesture, and typical password entry. Another advantage of face authentication is that it provides stronger defense against repudiation than token based authentication and password based authentication [28]. Besides face authentication, face identification is another application of face recognition, which compare a facial image with multiple registered users and identifies the user in the facial images. The face identification can cause privacy leakage in OSNs due to the identifiable personal images published in OSNs [3, 12]. Compared to their work, our study focuses on investigating the impact of the shared personal images that can be used to attack face authentication systems.

It is a well-known fact that face authentication is subject to spoofing attacks. An attacker can pass the authentication by displaying images or videos of a legitimate user in hard copy or on the screen [5]. But it is generally believed sufficiently secure as an authentication factor for common access protection, as an adversary usually has to be physically proximate to a victim in order to collected required face biometrics. Our findings indicate that this belief is not valid as the emergence of OSNFD. Face biometrics can now be disclosed in large scale and acquired by a remote adversary.

Liveness detection is the major countermeasure designed to mitigate the risk of spoofing attacks. Interaction based approach, multimodal based approach, and motion based approach are three popular types of liveness detection [29, 19, 4]. Interaction based approaches require real-time responses from claimants, including eye blink, head rotation, facial expression, etc. However, these approaches can be bypassed with one or two images [31]. Multimodal based approaches take face biometric and other biometrics into consideration together such as voice, facial thermogram, etc [29]. The multi-modal based approaches require additional hardware and specific environment. Motion based approaches are based on the detection of involuntary motions of a 3D face, such as involuntary rotation of head [19]. The approaches require high quality images captured with ideal lighting condition. Compared to these approaches, our estimation tool addresses this problem from a different perspective. Since OSNFD significantly compromise the confidentiality of face authentication, our tool is designed to increase the users' awareness before they publish their personal images so as to reduce the number of exploitable images available to an adversary.

7. CONCLUSION

In this paper, we investigated the threat of OSN-based facial disclosure (OSNFD) against some real-world face authentication systems. Our results show that the face authentication systems are vulnerable to OSNFD-based attacks. We analyzed the characteristics of these attacks from three major perspectives including security settings, target platforms and user behavior. The key attributes of the OSNFD were further extracted to develop a risk estimation tool that can help users understand the risks associated with their personal images shared in OSNs. Our work made the first step in systematically understanding the OSNFD. Quantitative evidence indicates that face authentication may not be a proper authentication factor as the confidentiality of face biometrics has been significantly compromised by OSNFD.

8. ACKNOWLEDGMENTS

The authors would like to thank Swee Won Lo (School of Information Systems, Singapore Management University) for her help and discussion on the user study and data collection. The authors would also like to thank Jianhui Huang (School of Information Systems, Singapore Management University) for his help and discussion on the the risk estimation model.

9. REFERENCES

[1] A. F. Abate, M. Nappi, D. Riccio, and G. Sabatino. 2d and 3d face recognition: A survey. *Pattern Recognition Letters*, 28(14):1885–1906, 2007.

[2] M. Abdel-Mottaleb and M. H. Mahoor. Assessment of blurring and facial expression effects on facial image recognition. In *Advances in Biometrics*, pages 12–18, 2005.

[3] A. Acquisti, R. Gross, and F. Stutzman. Faces of facebook: Privacy in the age of augmented reality. *BlackHat USA*, 2011.

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

[5] B. Biggio, Z. Akhtar, G. Fumera, G. Marcialis, and F. Roli. Security evaluation of biometric authentication systems under real spoofing attacks. *Biometrics, IET*, 1:11–24, 2012.

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

[7] A. Dantcheva, C. Chen, and A. Ross. Can facial cosmetics affect the matching accuracy of face recognition systems? In *Biometrics: Theory, Applications and Systems (BTAS), 2012 IEEE Fifth International Conference on*, pages 391–398, 2012.

[8] Facelock.mobi. http://www.facelock.mobi/facelock-for-apps.

[9] A. S. Georghiades, P. N. Belhumeur, and D. J. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 23(6):643–660, 2001.

[10] Google. http://www.android.com/about/ice-cream-sandwich/.

[11] Google. https://play.google.com/store/apps?hl=en.

[12] R. Gross and A. Acquisti. Information revelation and privacy in online social networks. In *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, pages 71–80, 2007.

[13] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker. Multi-pie. *Image and Vision Computing*, 28(5):807–813, 2010.

[14] A. J. Harris and D. C. Yen. Biometric authentication: assuring access to information. *Information Management & Computer Security*, 10(1):12–19, 2002.

[15] D. W. Hosmer Jr, S. Lemeshow, and R. X. Sturdivant. *Applied logistic regression*. Wiley. com, 2013.

[16] F. Hua, P. Johnson, N. Sazonova, P. Lopez-Meyer, and S. Schuckers. Impact of out-of-focus blur on face recognition performance based on modular transfer function. In *Biometrics (ICB), 2012 5th IAPR International Conference on*, pages 85–90, 2012.

[17] G. B. Huang, M. Mattar, T. Berg, E. Learned-Miller, et al. Labeled faces in the wild: A database forstudying face recognition in unconstrained environments. In *Workshop on Faces in'Real-Life'Images: Detection, Alignment, and Recognition*, 2008.

[18] A. K. Jain, A. Ross, and S. Pankanti. Biometrics: A tool for information security. *Trans. Info. For. Sec.*, 1(2):125–143, 2006.

[19] K. Kollreider, H. Fronthaler, and J. Bigun. Non-intrusive liveness detection by face images. *Image and Vision Computing*, 27(3):233–244, 2009.

[20] S. G. Kong, J. Heo, B. R. Abidi, J. Paik, and M. A. Abidi. Recent advances in visual and infrared face recognitionąła review. *Computer Vision and Image Understanding*, 97(1):103–135, 2005.

[21] N. Kumar, A. C. Berg, P. N. Belhumeur, and S. K. Nayar. Attribute and simile classifiers for face verification. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 365–372, 2009.

[22] I. Lab. https://www.idiap.ch/dataset/replayattack.

[23] K.-C. Lee, J. Ho, and D. J. Kriegman. Acquiring linear subspaces for face recognition under variable lighting. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 27(5):684–698, 2005.

[24] Lenovo. http://en.wikipedia.org/wiki/VeriFace.

[25] Luxand. http://www.luxand.com/.

[26] H. Moon and P. J. Phillips. The feret verification testing protocol for face recognition algorithms. In *Automatic Face and Gesture Recognition, 1998. Proceedings. Third IEEE International Conference on*, pages 48–53, 1998.

[27] E. Murphy-Chutorian and M. M. Trivedi. Head pose estimation in computer vision: A survey. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 31(4):607–626, 2009.

[28] L. O'Gorman. Comparing passwords, tokens, and biometrics for user authentication. *Proceedings of the IEEE*, 91(12):2021–2040, 2003.

[29] G. Pan, L. Sun, Z. Wu, and S. Lao. Eyeblink-based anti-spoofing in face recognition from a generic webcamera. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pages 1–8, 2007.

[30] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview of the face recognition grand challenge. In *Computer vision and pattern recognition, 2005. CVPR 2005. IEEE computer society conference on*, volume 1, pages 947–954, 2005.

[31] J. Rice. http://www.androidpolice.com/2012/08/03/android-jelly-beans-face-unlock-liveness-check-circumvented-with-simple-photo-editing/.

[32] C. J. V. Rijsbergen. *Information Retrieval*. Butterworth-Heinemann, 1979.

[33] SAS. http://www.sas.com/.

[34] T. Sim, S. Baker, and M. Bsat. The cmu pose, illumination, and expression database. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 25(12):1615–1618, 2003.

[35] F. Stutzman, R. Gross, and A. Acquisti. Silent listeners: The evolution of privacy and disclosure on facebook. *Journal of Privacy and Confidentiality*, 4(2):2, 2013.

[36] S. Trewin, C. Swart, L. Koved, J. Martino, K. Singh, and S. Ben-David. Biometric authentication on a mobile device: a study of user effort, error and task disruption. In *Proceedings of the 28th Annual Computer Security Applications Conference*, pages 159–168, 2012.

[37] VagueWare.com. http://www.vagueware.com/top-globally-popular-face-recognition-software/.

[38] Visidon. http://www.visidon.fi/en/Home.

[39] S. Vision. http://www.sensiblevision.com/en-us/home.aspx.

[40] K. Wagner. http://mashable.com/2013/09/16/facebook-photo-uploads/.

[41] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *Image Processing, IEEE Transactions on*, 13(4):600–612, 2004.

[42] J. D. Woodward. Biometrics: Privacy's foe or privacy's friend? *Proceedings of the IEEE*, 85(9):1480–1492, 1997.

[43] W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld. Face recognition: A literature survey. *ACM Computing Surveys (CSUR)*, 35(4):399–458, 2003.