Presentations and Attacks, and Spoofs, Oh My

Stephanie Schuckers, Clarkson University

sschucke@clarkson.edu

2/3/2016

Abstract

“Presentation attacks” are attacks at a biometric recognition data capture sensor which interfere with its normal operation. When artificial materials are used to create a fake biometric characteristic, it has been more commonly termed spoofing and the artefacts used to attack the system have been called spoofs. This paper gives an overview of this field, describes vocabulary formalized by the recent publication of an ISO standard for biometric “presentation attack detection”, and discusses evaluating the performance of systems which incorporate methods to detect and reject presentation attacks. In particular, this paper describes terms such as “attack presentation classification error rate” and “attack presentation match rate” from the ISO standard. To reduce confusion, it is important to differentiate metrics related to presentation attacks from the false accept rate which is specifically referring “zero-effort” imposter attempts, not a presentation attack. Last, some considerations for the future which relate to evaluating the performance of presentation attack detection are discussed.

Introduction

Attacks at the biometric sensor level, called ***presentation attacks***, are one vulnerability in biometric systems and are commonly known as “spoofing”. This year (2016), the first ISO standard has been published on biometric presentation attack detection [1]. This standard defines presentation attacks as the “presentation of an artefact or human characteristic to the biometric capture subsystem in a fashion that could interfere with the intended policy of the biometric system” [1]. One category of presentation attacks is the use of an artificial biometric or artefact to spoof a biometric device [2], e.g., Figure 1. Other examples include use of cadavers, modification of the biometric through surgery or mutilation, and non-conformant presentation.

Figure 1. Spoof fingerprint made of silicon (front) and associated mold used to form the fingerprint spoof (back).

Reports of spoof attacks in active identity management systems have included a South Korean woman in 2009 who used a special tape on her fingers to fool the fingerprint recognition system at a Japanese airport [3] and a Brazilian doctor who in 2013 used fake fingers made of silicone to sign in absent colleagues [4]. Additionally, attacks have also been highlighted by media groups, hackers, and universities. For example, the iPhone received attention when its new fingerprint reader located in the button of the phone was shown to be spoofed [5]. *While the threat is clearly there, there has been little evidence that widespread fraud based on spoofing has occurred in biometric systems.*

That being said, uses of biometrics for e-commerce and many other applications are growing and this type of fraud may yet gain importance; thus methods for mitigation are needed to reduce the risk of presentation attacks. Despite the fact that spoofing has been known for over a century, has been part of popular media for many decades, and has hundreds of academic literature going back almost twenty years, it is just recently that commercial solutions are regularly including language that suggests that robustness to presentation attacks is an important feature of their technology. (There were “early commercial adopters”, but these did not much gain much traction.) While the field has a history, confusion still abounds on what to call it, how much to worry about this type of fraud, how best to mitigate these attacks, and how to measure the effectiveness of those approaches. *The recently published ISO standards reflect international consensus on many of these topics and are an opportunity to begin to reduce the confusion.* This paper sets out to summarize the state of the art as well as some directions for the future.

Mitigation of Spoof Presentation Attacks

Mitigation can be accomplished using a combination of varied approaches. The use of multiple factors such as including components of what you have and/or what you know can reduce risk, as the attacker would need to both steal your biometric characteristic, create a successful spoof, and steal the other factors (e.g. steal a password/pin). For example, the FIDO Alliance [6] is an emerging industry standard for authentication. Some implementations incorporate a combination of a locally stored biometric for user verification and asymmetric key cryptography for device authentication. In this implementation, an attacker would need to steal both a user’s biometric and a user’s mobile device. This is difficult to mount in a large scalable way. Multiple biometrics can also be used; however, the way the information is fused is important if it is being used to make it difficult to spoof the system. Other examples of mitigation include limits on the number of biometric attempts since creating a successful spoof can be difficult. Use of supervision can also be useful for some applications (e.g. border security); however, those supervising the process should be trained on how to look for spoof attacks, since some attacks can be quite surreptitious. For example, a fingerprint spoof can be made very thin, skin colored, and glued to an attacker’s real finger in order to not be noticed (Figure 2). A challenge-response can be effective to repel some types of attacks. Examples include asking the user to say a specific phrase, blink, or make a facial expression. For some applications, a mechanism to report device loss and the ability to delete, revoke, or not accept the credential from the stolen device. This limits the time available to mount a successful spoof attempt. Last, detection of the presentation attack itself through methods, such as liveness detection or artefact detection, can be used and are known in the ISO standard as ***presentation attack detection***. More detail is provided in the next sections.

Figure 2. Biometric spoofs can be created which is difficult to spot, e.g. a skin colored, thin fingerprint spoof can be glued to a user's actual finger.

In summary, the following approaches can be used for mitigation of spoof presentation attacks:

- Multi-factor
- Multi-biometrics
- Limit number of attempts or timeout mechanisms
- Supervision with appropriate training
- Challenge-response
- Device loss procedures
- Presentation attack detection

While presentation attacks are a vulnerability to be aware of, *it is important to perform an analysis of the specific application which uses biometric recognition when deciding whether this type of fraud has significant risk and to determine the best ways to mitigate which are balanced in terms of the risk/cost tradeoff.*

Use of stolen fingerprint databases to mount attacks

There is quite a bit of confusion regarding stolen biometric information, such as the stolen database of fingerprint images in the US Office of Personnel Management hack of 2015 [7]. It is quite worrisome that hackers may have access to an individual's fingerprint and, like other private data that is stolen, the uses of this data can be quite damaging to an individual. However, there is a misconception that a stolen fingerprint is the equivalent of a stolen password. The difference relates to the input mechanism. A password can be inputted quite simply by entering the characters through any keyboard. A biometric in concept needs to be entered through a biometric capture device. Biometric images are not entered directly. Instead, the stolen images would need to be converted into a spoof artefact which can be used in order to be measurable by the data capture sub-system. With many biometric implementations which have local storage of the enrolled template, the stolen artefact is only useful if the device with the associated stored enrollment is present. In another attack type with a stolen fingerprint database, a large set of biometric spoofs could be created and a brute force attack be

attempted using these many physical artefacts. However, it would require a significant undertaking to create this sort of spoof library as well as a significant amount of time to carry out a brute force attack using the many created spoofs.

Outside of presentation attacks, the stolen fingerprint image can only be used directly (i.e. without creating a physical spoof) by bypassing the biometric data capture device, i.e., a fingerprint reader, and inserting it prior to the feature extraction software. For this attack to be successful, the security of the connection would need to be broken. In fingerprint readers in mobile device, this is often a secured connection to a protected part of the device where the fingerprint matching occurs. In other applications, the communication of the fingerprint image is (and should be) secured through other means. The raw fingerprint image is simply not useful, as is, at any other point in the biometric process. To be used later in the process, the raw fingerprint image would need to be converted to a feature set, i.e., the subset of the biometric image used for recognition. Each biometric system has a specialized features and a specific format of a template. This requires significant knowledge of the system under attack which for most consumer applications is a trade secret and associated malware to interfere in the biometric matcher software itself. If malware is injected, it is more likely that other means of corrupting the software would be easier to employ, e.g., overriding the matching decision. *In summary, while these approaches to use a stolen fingerprint database for an attack are possible, they are not as simple to mount as finding a keyboard and entering the stolen password.*

Presentation Attack Detection (or more commonly known as liveness detection)

Presentation attack detection (PAD) methods recognize a presentation attack. ISO standards define it as “automated determination of a presentation attack” [1]. *The goal of presentation attack detection is to determine if the biometric being captured is an actual measurement from the authorized, live person who is present at the time of capture* [8], i.e., a “normal presentation”. Examples include hardware-based sensors which are built into fingerprint scanners to measure life signs and software methods which analyze the fingerprint image for evidence of spoofing. Other names for PAD and can be considered as sub-sets of PAD include liveness detection, artefact detection, spoof detection, or anti-spoofing, among others.

Additionally, a biometric-based challenge-response combined with PAD systems can be used as a method for mitigating presentation attacks and could be compared to the CAPCHA or similar methods to protect from bots when a password is entered. A biometric-based challenge response can be voluntary (e.g. say a specific phrase that is displayed or blink) or involuntary (pupil dilating in response to specific illumination pattern).

An advantage of implementing PAD is that this addresses two common concerns related to “biometrics are not secret” and “biometrics cannot be changed”. For a biometric system which rejects presentation attacks and is secure throughout the system, the fact that the biometric is not secret or cannot be changed is less important as the system recognizes when someone is trying to attack with the stolen biometric information. Methods to ensure a template is “cancellable” can also be used.

Performance Evaluation of PAD

The assessment of the effectiveness of these systems is critical and is aided by biometric standards, e.g. ISO/IEC SC 37, WD 30107 [1], and benchmark competitions for liveness detection, e.g. LivDet which has been held biennially since 2009 [9]. Additionally, many research papers have focused on the effectiveness of the PAD module alone [e.g., 8, 10, 11] and in combination with the matcher [e.g., 12, 13].

This following section summarizes various testing modes for PAD and metrics associated with these modes. The metrics are also emerging from ISO standards ISO/IEC SC 37, WD 30107 Part 3 [14]. While many papers have studied this area, the metric titles have varied as might happen with any new emerging field. Consistently using the international standards would be very helpful as PAD solutions enter the marketplace. It is quite common to hear people talk of the same things with different names. Examples of metrics that have been used for this have been called: classification rate (percent correctly classified), ferrfake (rate of misclassified fake biometrics), spoof false accept, among others. Many have used the term false accept rate (FAR). We suggest that FAR is too easily confused with matching performance and should **not** be used. Further detail is provided in the next paragraph.

PAD assessment SHOULD NOT be measured by the False Accept Rate

There has been quite a bit of confusion around metrics related to spoofing when calling a successful presentation a false accept and measuring its performance using the false accept rate. According to ISO, a false accept rate is the “proportion of verification transactions with wrongful claims of identity that are incorrectly confirmed” [15]. As further defined, “claims of identity” are directly associated with a false match rate. A false match rate is “proportion of zero effort imposter attempt samples falsely declared to match the compared non-self template.” The key word in this definition is “zero-effort”. This refers to imposter attempts that use their own biometric. This metric is specifically related between individual comparisons of one individual to another or “inter-class” distribution. Zero-effort imposter testing is typically performed comparing biometric samples between different individuals who are necessarily “live”. Spoofing or “non-zero effort” attacks are quite another matter. In this, a “spoof” or “artefact” needs to be created from a stolen or generated biometric characteristic. To test the system’s robustness against these types of attacks, a set of spoof attacks needs to be created and used in testing. It can easily be seen that this is quite a different test than “inter-individual” comparisons to compute FAR. The metrics to quantify performance related to presentation attacks will be described in detail below and include “Attack Presentation Classification Error Rate” and “Attack Presentation Match Rate”.

To recap:

- Zero-effort imposter testing—“inter-individual” testing is used to estimate false match rate and false accept rate (FMR and FAR)
- Non-zero effort testing—spoofs or artefacts are created to test system robustness to recognizing presentation attacks and are used to estimate spoof performance metrics summarized below (APCER, APMR)

Performance Metrics for Evaluation of PAD

This section was also outlined in an abstract as part of the IBPC 2014 [16]. The metric names themselves have evolved from that time in the developing ISO standards and are edited here to reflect those changes.

Performance metrics for PAD module alone

For assessing a PAD module, there are two testing modes:

(A) Presentation Attack: Performance when tested with a presentation attack, e.g., presenting a fake fingerprint made of gelatin.

(B) Normal Presentation: Performance when tested in normal operation (not a presentation attack), e.g. presenting a live finger.

Associated metrics are:

Attack Presentation Classification Error Rate (APCER): The PAD module alone is assessed by performing a presentation attack repeatedly and measuring the response of the system. An example would be test a fingerprint PAD module with a set of spoof fingerprints.

Normal Presentation Classification Error Rate (NPCER): The PAD module alone is assessed by repeatedly using the module in normal operation. An example would be testing a fingerprint PAD module with a group of live subjects with normal fingerprints (i.e. not a presentation attack). *No PAD test should be performed without also considering NPCER as it is a trade-off with APCER.*

An example of results from this type of evaluation is shown from the LivDet competitions in Figure 3.

Figure 3. Average of APCER and NPCER for top two performers across all databases tested in LivDet competitions hosted 2009, 2011, 2013, 2015.

Performance metrics for PAD module and matcher combined

When PAD module is combined with a matcher, there are three testing modes:

(A) Presentation Attack (with an attack made using a genuine characteristics)

(B) Normal Presentation—Genuine: an individual matched with their own reference template

(C) Normal Presentation—Imposter: an individual matched with another's reference template

(Note: It is possible to test a Presentation Attack with an Imposter. This is not typical and will not be covered here.)

Associated metrics are as follows:

(A) Testing Mode: Presentation Attack—Genuine

Attack Presentation Match Rate (APMR): When testing presentation attacks, an attack that both passes a PAD module and matches the stored biometric (the case from above: Normal Presentation—Match) is defined as the Attack Presentation Match Rate. Otherwise the presentation attack is correctly detected by either the PAD module or the matcher or both.

(B) Testing Mode: Normal Presentation—Genuine

False Reject Rate (FRR). For Genuine cases, the correct decision is Normal Presentation—Match.

Attempts which are called a Presentation Attack OR a Non-Match are called a *False Reject*. The associated metric is the false reject rate. Note that the PAD module has an impact on the FRR.

(C) Testing Mode: Normal Presentation—Imposter

False Accept Rate (FAR). For Imposter cases, the correct decision is Normal Presentation—Non-Match. Attempts which are called Normal Presentation—Match are considered a *False Accept*. The associated metric is the false accept rate. This has also been called the “zero effort” false accept rate, indicating that this only includes imposter, i.e., non-matching subjects, and not those that make an effort to attack the system.

Traditionally tradeoff curves (like ROC and DET) plot FRR and FAR as a function of matching threshold. For the combination of PAD module and matcher, a 3D assessment space is created with a trade-off between APMR, FRR, and FAR. Since the PAD module may reject a Normal Presentation, it has an impact on BOTH the FRR and APMR. Further complexity is added because there are potentially at least two decision thresholds: one for the PAD module and one for the matcher. In [16], using public dataset, a more detailed example is presented focusing on each of these metrics, as well as example trade-off curves.

The Future

In summary, the ISO standards have been significant to developing a common terminology, testing framework, and performance metrics which the biometrics community and application architects can use. So what is next? There is a need to triage the types of attacks to develop solutions for and to test robustness against, in other words, separate more simple attacks from those that are more sophisticated. Attacks can be separated into levels based on an increasing level of difficulty to mount based on time, skill, and equipment, based on frameworks such as Common Criteria [17]. In this paper, Table 1 is a strawman which separates presentation attacks by level and modality. Level A attacks are attacks which are quite simple to carry out and require relatively little time, expertise, or equipment. An additional aspect of Level A is that the biometric characteristic under attack is quite easy to obtain (e.g. face image from social media, fingerprint from the device and reused directly). Level B attacks require more time, expertise and equipment. Additionally, the difficulty to acquire the biometric characteristic is higher (e.g., a latent fingerprint converted to 3D spoof, high quality video of a person’s face). Level C includes the most difficult attacks. Each level could potentially have corresponding increased testing requirements such that systems are robust against attacks for that level as well as the levels below it.

As the maturity of this field increases and as more commercial implementation emerge, a robust testing framework that goes beyond a checkbox mentality is needed (e.g., Do you have liveness? Yes or No.). It is hoped that the inclusion of performance results related presentation attacks will become as common place as our reporting (and understanding) of false reject rate and false accept rates. Complicating matters, detailed information will need to be provided on what types of attacks a specific solution is robust, in the midst of a changing attack space.

Table 1. Spoof presentation attacks separated by levels based on time, expertise, and equipment

		Fingerprint	Face	Iris	Voice
Level A	Time: short Expertise: anyone Equipment: readily available	paper printout, direct use of latent print on the scanner	paper printout of face image, mobile phone display of face photo	paper printout of iris image, mobile phone display of iris photo	replay of audio recording

	Source of biometric characteristic: easy to obtain	lift of fingerprint	photo from social media	photo from social media	recording of voice
Level B	Time: >3 days Expertise: moderate skill and practice needed Equipment: available but requires planning	fingerprints made from artificial materials such as gelatin, silicon.	paper masks, video display of face (with movement and blinking)	video display of an iris (with movement and blinking)	replay of audio recording of specific passphrase, voice mimicry
	Source of biometric characteristic: more difficult to obtain	latent print, stolen fingerprint image	video of subject, high quality photo	video of subject, high quality photo	recording of voice of specific phrase
Level C	Time: >10 days Expertise: extensive skill and practice needed Equipment: specialized and not readily available	3D printed spoofs	silicon masks, theatrical masks,	contacts lens with a specific pattern	voice synthesizer
	Source of biometric characteristic: more difficult to obtain	3D fingerprint information from subject	3D face information from subject	high quality photo in Near IR	multiple recordings of voice to train synthesizer

Acknowledgements: This material is based upon work supported by the Center for Identification Technology Research and the National Science Foundation under Grant #1068055 and National Institute of Standards and Technology (NIST) #60NANB10D216.

¹ International Standards Organization, ISO/IEC TC JTC1/SC 37, ISO/IEC 30107-1:2016, Information technology -- Biometric presentation attack detection -- Part 1: Framework, 1/15/2016.

² T. Matsumoto, H. Matsumoto, K. Yamada, and S. Hoshino, "Impact of Artificial 'Gummy' Fingers on Fingerprint Systems," in *Proceedings of SPIE Vol. 4677*, 2002, vol. 4677.

³ "AFP: SKorean fools finger printing system at Japan airport: reports." [Online]. Available: <http://www.google.com/hostednews/afp/article/ALeqM5jwMI9y-RtlCG0LXfKIF5yX0uxgzg>. [Accessed: 17-May-2013]

⁴ "Doctor 'used silicone fingers' to sign in for colleagues," *BBC News*, 12-Mar-2013. [Online]. Available: <http://www.bbc.co.uk/news/world-latin-america-21756709>.

⁵ "Chaos Computer Club breaks Apple TouchID." [Online]. Available: <http://www.ccc.de/en/updates/2013/ccc-breaks-apple-touchid>

⁶ The FIDO Alliance, <https://fidoalliance.org/>

⁷ A. Peterson, "OPM says 5.6 million fingerprints stolen in cyberattack, five times as many as previously thought" *The Washington Post*, 9/23/2015.

⁸ R. Derakhshani, S. A. C. Schuckers, L. A. Hornak, and L. O'Gorman, "Determination of vitality from a non-invasive biomedical measurement for use in fingerprint scanners," *Pattern Recognition*, vol. 36, no. 2, pp. 383–396, Feb. 2003.

⁹ Liveness Detection (LivDet 2009, 2011, 2013, 2015) Competitions and Datasets, <http://www.clarkson.edu/biosal/Liveness%20Resources.html>

¹⁰ Y. Zhang, J. Tian, X. Chen, X. Yang, and P. Shi, "Fake Finger Detection Based on Thin-Plate Spline Distortion Model," *Advances in Biometrics*, pp. 742-749, 2007.

-
- ¹¹ Z. Wei, X. Qiu, Z. Sun, and T. Tan, "Counterfeit iris detection based on texture analysis," *19th International Conference on Pattern Recognition ICPR 2008*, pp. 1-4, 2008.
- ¹² E. Marasco, P. Johnson, C. Sansone, and S. Schuckers, "Increase the security of multibiometric systems by incorporating a spoof detection algorithm in the fusion mechanism," *Multiple Classifier Systems*, pp. 309-318, 2011.
- ¹³ A. Anjos, I. Chingovska, and S. Marcel, "Anti-spoofing in action: joint operation with a verification system," *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, Workshop on Biometrics*, 2013.
- ¹⁴ International Standards Organization, ISO/IEC TC JTC1/SC 37, ISO/IEC 30107-3, Information technology -- Biometric presentation attack detection -- Part 3: Testing and reporting, Under Development.
- ¹⁵ International Standards Organization, ISO/IEC 19795-1:2006 Information technology -- Biometric performance testing and reporting -- Part 1: Principles and framework.
- ¹⁶ P Johnson and S Schuckers, "Evaluation of presentation attack detection", International Biometric Performance Conference, 2014.
- ¹⁷ O Henniger, D Scheuermann, and T Kniess On security evaluation of fingerprint recognition systems, IBPC, 2010, http://biometrics.nist.gov/cs_links/ibpc2010/pdfs/Henniger2_Olaf_IBPC_Paper.pdf