Robustness of Multi-modal Biometric Systems under Realistic Spoof Attacks against All Traits

Zahid Akhtar, Battista Biggio, Giorgio Fumera, and Gian Luca Marcialis
Department of Electrical and Electronical Engineering, University of Cagliari
Piazza d'Armi, 09123 Cagliari, Italy
E-mail: {z.momin,battista.biggio,fumera,marcialis}@diee.unica.it
URL: http://prag.diee.unica.it

***Abstract*—Spoof attacks consist in submitting fake biometric traits to biometric systems, and are a major threat that can curtail their security. Multi-modal biometric systems are commonly believed to be intrinsically more robust to spoof attacks, but recent works have shown that they can be evaded by spoofing even a single biometric trait. This result was however obtained under the worst-case scenario that the attacker is able to fabricate an exact replica of the genuine biometric trait, which was simulated by assuming that the matching score distribution of fake traits is identical to the one of genuine users. This demands for a more thorough investigation of the robustness of multi-modal biometric systems against *realistic* spoof attacks, namely under non-worst case scenarios. In this paper we focus on bi-modal systems made up of a face and a fingerprint matcher, whose scores are fused using the well-known sum, product, weighted sum and likelihood ratio (LLR) rules. We evaluate their robustness against *realistic* spoof attacks obtained by fabricating fake biometric traits. The main goal of our study is to investigate whether a realistic spoof attack against *both* modalities can allow the attacker to crack the multimodal system. Our results show that even in a realistic, non-worst case scenario, the false acceptance rate (FAR) remarkably increases.**

I. Introduction

The forgery in unattended operation of biometric authentication has escalated with the rapid growth in the use of biometric systems. Several researchers have shown that many biometrics such as face, fingerprint and iris of legitimate users can be stealthy procured to produce synthetic (fake) biometric traits to attack biometric sensors [1], [2], [3]. This kind of biometric forgery is know as "spoof attack". Since spoof attacks are carried out directly on the biometric sensor, they are also named as "direct attacks".

Multi-modal biometric systems are commonly believed to be intrinsically robust against spoof attacks, contrary to mono-modal systems [4]. Multi-modal systems integrate evidences produced by multiple source of information, namely different biometric traits, and have been originally proposed to improve accuracy and to overcome some inherent limitations of mono-modal systems. The belief on their intrinsic robustness against spoof attacks is based on the assumption that an intruder has to spoof *all* fused biometrics *simultaneously* to crack the system. However, this assumption is not based than on any theoretical or empirical evidence.

Contrary to the above common belief, recent works have shown that multimodal biometric systems can be cracked by faking even a *single* biometric trait [5], [6], [7]. However, most of these results were obtained under the worst-case scenario that the attacker is able to fabricate a perfect replica of the genuine biometric trait, which was *simulated* by assuming that the distribution of matching scores of fake traits is identical to the one of genuine users. Accordingly, these results raise the issue of investigating more thoroughly the robustness of multi-modal biometric systems against *realistic* spoof attacks, namely under non-worst case scenarios.

In our previous works we addressed this issue by proposing a model of the matching score distribution of fake traits to simulate non-worst case scenarios [8], [9], and by analysing to what extent the non-worst case assumption is representative of realistic attacks [10]. In this paper, we focus instead on verifying the common belief mentioned above, that a multi-modal system can actually be cracked by spoofing *all* the considered biometric traits. To this aim, we empirically evaluate the robustness of bi-modal systems made up of a face and a fingerprint matcher, by fabricating fake fingerprints and fake face images using realistic techniques. We point our that [6] and [7] considered spoof attacks against more than one matcher, but one or both attacks were simulated under the worst-case assumption. We carry out our empirical investigation using several well-known score fusion rules: sum, product, weighted sum and likelihood ratio (LLR). Our results show that even in a realistic, non-worst case scenario, the false acceptance rate (FAR) can remarkably increase when both traits are spoofed, regardless on the score fusion rule. This provides a clear evidence that multi-modal systems can be cracked through *realistic* spoof attacks, at least when *all* biometric traits are spoofed.

The paper is organized as follows. Section II provide a concise description of the architecture of multi-modal biometric systems, with reference to the one considered in this paper, and summarises previous works their robustness against spoof attacks. The data sets used in this work are described in Section III. The experimental results are reported and discussed in Section IV. Section V concludes the paper.

II. Background

A. Multi-modal Biometric Systems

Figure 1 illustrates the architecture of a multi-modal biometric system, with reference to the one considered in this

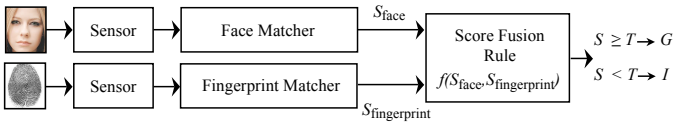

Fig. 1. Outline of a multi-modal system composed of a face and a fingerprint matcher, whose matching scores are combined by a fusion rule $S = f(S_{\text{face}}, S_{\text{fingerprint}})$.

work, namely a bi-modal system composed of a face and a fingerprint matcher. In a multi-modal systems, information fusion can be carried at various levels: feature extraction level, matching score level and decision level. Fusion at the matching score level is generally preferred due to ease in accessing and combining of matching scores, and will be adopted in this work, as in [5], [6], [7].

In a *verification* setting, first the user provides his face and fingerprint to the corresponding sensors, and claims his identity. Then, the biometric traits are individually processed, and compared with the enrolled respective templates of the claimed identity. Lastly, the matching scores outputted by each matcher (S_{face} and $S_{\text{fingerprint}}$ respectively for the face and fingerprint matchers) are fused, according to a given rule. If the fused score S is equal or greater than a predefined threshold T, then user is accepted as genuine (this decision is denoted as 'G' in the figure), otherwise it is rejected as an impostor ('I').

B. Multi-modal Biometric Systems under Spoof Attacks

Spoof attacks against biometric systems have been investigated since the early paper of [2], where a 60% acceptance rate was reported for artificial fingerprints replicated using gum and gelatin. Some researchers have also shown that iris may be faked by suitable contact lenses [3], and face verification systems can be fooled by showing a simple photograph of a genuine user [13]. Several liveness/vitality detection methods have been proposed by most of the researchers as a possible countermeasure. For instance, in fingerprint liveness detection many hardware- and software-based techniques have been developed and adopted in commercial products (e.g., see the survey in [14]). However, liveness detection has a major disadvantage: it typically increases the percentage of genuine users rejected as impostor (i.e., the false rejection rate, FRR). Other spoof detection algorithms have been proposed for other biometric traits, but their performance was not satisfactory as well.

Besides liveness detection, as explained in Section I multi-modal biometric systems are commonly believed to be themselves a "natural" defence mechanism against spoof attacks, based on the intuitive assumption that they require spoofing of *all* fused biometric traits simultaneously in order to be cracked. However, there is no theoretical or empirical evidence to support this belief. On the contrary, some recent works [5], [6], [7] have addressed this issue and questioned the validity of the above claim. Reported results in [5] and [6] using sum, weighted sum and LLR score fusion rules showed that multimodal systems can be cracked by spoofing *only one* biometric trait. They considered systems made up of two (face and fingerprint) or three matchers (face, fingerprint and iris). A substantial increase of the false acceptance rate (FAR) of these systems under spoof attacks was indeed highlighted. However, most of these results were obtained using simulated spoof attacks at the matching score level, under the worst-case assumption that the attacker is able to fabricate the exact replica of the genuine biometric traits, and thus that the matching score distribution of the spoofed biometric traits is equal to the one of the genuine traits. While this worst-case assumption may be reasonable for some biometric traits, like 2D faces (as discussed in [7]), its validity for other traits like iris and fingerprint is questionable.

Notably, in [7] some experiments have been carried out using a fraction of real spoofed fingerprints from the Fingerprint Liveness Detection Competition (*LivDet09*) [16]. The fake fingerprint score distribution obtained in the study was remarkably different from the genuine score distribution, additionally providing the proof that worst-case hypothesis is not realistic for all biometric traits and spoofing techniques.

The above results raise the issue of more thoroughly investigate the robustness of multi-modal systems against *realistic* spoof attacks, namely their performance degradation in the non-worst case scenario when the fake traits are not exact replicas of genuine ones. This is however a difficult task, as in principle it requires to fabricate fake traits. In our previous works we started to addressed this issue, under two different viewpoints. On the one hand, we proposed a model of the matching score distribution of fake traits to *simulate* also non-worst case scenarios [8], [9], thus extending the scope of the analysis of [5], [6], [7], without the need of fabricating fake traits. On the other hand, we fabricated fake fingerprints and fake face images to assess to what extent the non-worst case assumption is representative of realistic attacks against these two traits [10]. As explained in Section I, in this paper we focus on one of the other issues related to the robustness of multi-modal systems against spoof attacks, namely on verifying the common belief that they can be cracked by spoofing *all* the considered biometric traits.

III. Experimental setup

In this study, we used a bi-modal face-fingerprint system like the one of Fig. 1, with different fusion rules. To evaluate the performance of the system when both traits are spoofed, we constructed a data set of faces and fingerprints, and the fabricated fake traits for both modalities.

A. Fusion rules

We used the following fixed and trained fusion rules. Note that in fixed fusion rules no parameters have to be set during the design of the system, contrary to trained rules.

1) **Fixed rules**

Sum. The fused score is obtained by simple addition of the individual score values:

$$S = S_{\text{face}} + S_{\text{fingerprint}} \,. \tag{1}$$

Product. The product rule computes the fused score as:

$$S = S_{\text{face}} \times S_{\text{fingerprint}} . \quad (2)$$

2) **Trained rules**

Weighted sum by Linear Discriminant Analysis (LDA). The individual scores are linearly combined as:

$$S = w_0 + w_1 S_{\text{face}} + w_2 S_{\text{fingerprint}} , \quad (3)$$

where w_0, w_1 and w_2 are weights determined by the maximization of the Fisher Distance (FD) between the score distributions of genuine and impostor users. Let μ_I, σ_I^2 and μ_G, σ_G^2 be the mean and variance of impostor ('I') and genuine ('G') score distributions, respectively. In the case of two matchers, FD is given by:

$$FD = \frac{(\mu_G - \mu_I)^2}{\sigma_G^2 + \sigma_I^2} . \quad (4)$$

Likelihood ratio (LLR). This rule computes the fused score as follows, when s_{face} and $s_{\text{fingerprint}}$ are considered as independent random variables:

$$S = \frac{p(S_{\text{face}}|G) \cdot p(S_{\text{fingerprint}}|G)}{p(S_{\text{face}}|I) \cdot p(S_{\text{fingerprint}}|I)} , \quad (5)$$

where p(·|G) and p(·|I) are the matching scores probability density function (PDF) of genuine and impostor users, respectively. In general, parametric (e.g., Gaussian, Gamma, Beta) or non-parametric models (e.g., Parzen windows) can be used to fit the genuine and impostor distributions.

B. Data set

We used two separate data sets of faces and fingerprints previously collected by our research group in different times, coming from different persons. To build a multi-modal data set (i.e., a data set in which each user has a face and a fingerprint trait), since the two data sets did not contain the same users, we randomly and uniquely combined the face modality and the fingerprint modality of pairs of *different* users of respective data sets, obtaining a "chimerical" data set of 40 users, with 40 genuine samples per user. This is a common procedure used in works on multi-modal systems [19]. To carry out more runs of the experiments, we repeated the above process for five times, obtaining five different chimerical data sets. For each data set, we used 40% of the users as training set, and the remaining ones as testing set. We estimated the parameters of the trained score fusion rules on the training set. The presented results are average testing set results on the five data sets.

The fingerprint and the face verification systems used for the experiments were implemented using the NIST Bozorth3 matching algorithm [20] and the EBGM matching algorithm [21], respectively.

In weighted sum rule, the weights were computed on training set while in LLR rule, we fitted a parametric distribution to training set for individual density estimation. We opted Gamma distribution, as in [5], because it was providing a better approximation of our data.

Fig. 2. Example of face spoof. Original template image (left), spoofed face by "photo-attack" (right).

C. Spoofing techniques and spoof data sets.

Starting from the data sets described above, we constructed spoof attacks as follows.

1) **Face Spoofing**

Spoofing techniques clearly depend on the nature of the system that has to be attacked. In the case of face spoofing, one has to consider that face recognition systems can be widely categorized into two classes: 2D (2-dimensional) and 3D (3-dimensional) systems. A biometric 2D face recognition system takes into consideration only the two dimensional image of the face. 3D systems are clearly more complex, and recognize faces on the basis of features extracted from the 3D shape of the whole face, using methods such as paraxial viewing, or patterned illumination light [17].

Latest studies on the vulnerability of face recognition systems to attacks have shown that they are widely susceptible to spoof attempts. The 2D face recognition systems can be easily misled by presenting a photograph or a facial video of the enrolled person's face displayed on a portable device [1], [18]. A genuine user's facial image or video can be simply captured using small cameras, without their awareness and consent for spoofing purpose. Moreover, in today's modern network society, person's facial photographs are usually available to the public. For example, to fool a biometric authentication system an attacker can submit a genuine user's photograph acquired from a social networking web site. Face video spoofing of genuine users has behavioral clues like expressions and eye blinking, which may also thwart vitality detection techniques.

3D face recognition systems are spoofed by 3D face model or face mask fabricated by rubber, plastic, silica gel [18]. The prerequisite to spoof 3D face systems is the presented face to camera should be three dimensional, thus making 3D face spoofing more complicated than a 2D face spoofing. Due to ease in spoof fabrication, and to the wide adoption of 2D systems, photographs and videos of faces are still the most common methods for spoofing faces.

To date, no face biometric data set with spoof attack samples have been publicly released. Hence, we fabricated spoofed face images using the face data set previously described. Live faces in our data set were collected with different lighting and

Fig. 3. Example of fingerprint spoof. Original template image (left), spoofed fingerprint by silicon (right).

facial expressions conditions, with a time interval of about two weeks between the two sessions. We fabricated the spoofed faces using the "photo attack" method mentioned in [1], [18]. It consists in putting in front of the camera the photograph of each genuine user. We show the photograph to the camera using a laptop screen. Figure 2 depicts an example of a live and spoofed face of a genuine user, fabricated by "photo attack", acquired by the camera. No visible artifact arises from the use of this technique, at least in this example. Namely, the spoofed face looks exactly like a genuine one.

2) **Fingerprint Spoofing**

Spoofed fingerprint can be replicated by two methods: "consensual" or "direct casts", and "non-consensual" or "indirect casts". When the consensual method is used, fake fingerprints are fabricated by user's consent and cooperation. The non-consensual method, instead, does not require any cooperation of the user, since latent finger-marks are used to fabricate the spoofed fingerprint. Notably, most of the researchers reported results on spoofed fingerprints reproduced by the consensual method [2], [16].

In the consensual method users press their finger in a suitable soft material such as wax, play doh, dental mold material, or plaster, to create the negative impression as a mold. From mold fake fingerprints are casted using different materials like silicon, moldable plastic, plaster, or clay.

For spoofing the fingerprints, we adopted the "consensual" method. We used a plasticine-like material as the mold while a two-compound mixture of liquid silicon and a catalyst were used as cast. The main property of the cast material is high flexibility silicon resin (SILGUM HF) with a very low linear shrinkage. All live and spoofed fingerprint images were collected using the Biometrika FX2000 optical sensor. Figure 3 shows the original, "live" client image, beside a replica made up of silicon.

IV. Experimental Results

In this section, we evaluate whether realistic spoofing attacks involving all the modalities of the considered multi-modal systems, using the four score fusion rules mentioned above, allow the attacker to break them. For reference, we also report the results under the worst-case assumption. The results are reported in Figure 4 in terms of the average Detection Error Trade-off (DET) curves attained on the test set. A DET curve reports the false rejection rate (FRR) as a function of the FAR, both computed parametrically being equally the decision threshold T on the fused score. Note that the FAR under a spoof attack is defined as the percentage of spoof attempts that got accepted as genuine, which is also referred to as Spoofed FAR (SFAR) in [6].

In Figure 4 we show the DET curves of the mono- and the multi-modal systems under normal operation (i.e., with no spoofing attack), using solid curves, and the performance under a spoof attack against one trait (both for the mono-modal and multi-modal-systems) and both traits (for the multi-modal-systems), using dashed curves. For reference, we also report the DET curve corresponding to a worst-case attack against both traits, which was *simulated* as in [5], [6], [7]. Note that the latter DET curve corresponds to the line $FAR = FRR$, as the matching score distribution of fake traits is assumed to be identical to the one of genuine users. Note also that the red (fingerprint individual system) and green curves (face individual system), as well as the black dashed curve (worst-case attack against both traits of the multi-modal system) are the same in all plots. Therefore, only the two blue curves corresponding to multi-modal systems without spoofing attacks, and with realistic spoofing attacks against both traits, change depending on the fusion rule.

We can first observe that, under normal operation (i.e., no spoofing attacks), fusion improved the performance of the corresponding mono-modal systems. The only exception are the sum and weighted sum rules, at high FAR vales. This behavior is due to the fact that the genuine and impostor score distributions of the face matcher in the considered data set turned out to be more overlapping than the ones produced by the fingerprint matcher, as it can be noted from the worse DET curves obtained by the individual face system with respect to the individual fingerprint system. In other words, generally the benefit of fusion are exploited when the classifiers show complementary nature. However, since in this case one of the modality is significantly worse over all the DET curve, the performance of multimodal system is below the best performing classifier.

Consider now the DET curves related to the realistic spoof attacks against mono- and multi-modal systems. It can be seen that the performance under attack considerably worsen both for the mono-modal face and fingerprint system. The performance of the multi-modal system under a realistic spoof attack against *both* traits is however better than the one attained by both mono-modal systems under attack, for all score fusion rules. In other words, an attacker has lower chances to evade the bi-modal systems considered in our experiments, when he spoofs both traits, than to evade each single mono-modal system. Accordingly, we can say that the bi-modal systems considered in our experiments exhibited a higher robustness to *realistic* spoof attacks against both traits, than the corresponding mono-modal systems against the same attacks.

This provides evidence to the common claim that multi-modal systems are more robust to spoofing attacks than mono-modal ones.

However, a comparison of the solid and dashed blue curves clearly shows that the performance of the multi-modal systems under a *realistic* spoofing attack against *both* traits is significantly worse than under normal operation, for all the fusion rules considered. This indicates that the probability of an impostor evading the multi-modal systems is high, even if he does not fabricate a perfect replica of the spoofed trait. For instance, using the product rule, if the ZeroFAR operating point is chosen on the training set (namely, the lowest decision threshold T which provides a zero FAR on training samples), an average FAR of 0.51% is obtained on testing samples under normal operation. When both face and fingerprints are spoofed instead, the FAR increases to 27.01%.

Accordingly, as far as our results are concerned, they provide evidence to the common claim that multi-modal systems can be cracked by spoofing attacks against *all* matchers, even when the attacker is not able to fabricate the exact replica of genuine user's biometric traits. Finally, note that the performance attained by the multi-modal systems under such *realistic* spoofing attacks is much better (although still not suitable for the requirements of security applications) than the one predicted under the worst-case assumption (black dashed line). This provides further evidence, besides the one of [10], that the worst-case assumption can be not representative of real spoofing attacks. This also further motivates the need of novel methods to evaluate the robustness of biometric systems under spoofing attacks, without requiring the fabrication of fake traits, as in [8], [9].

V. Conclusion and future work

Our experiments on *realistic* spoofing attacks provided evidence of two common beliefs about the robustness of multi-modal biometric systems. First, they can be more robust than each corresponding mono-modal system, even in the case when *all* biometric traits are spoofed. Second, their performance under a spoofing attack against all traits is still unacceptable for security applications. In other words, they can be cracked by spoofing *all* the fused traits, even when the attacker is not able to fabricate an exact replica of the genuine user's traits. We also found that the worst-case assumption considered in previous works [5], [6], [7] can be not representative of realistic spoofing attacks: its suitability may depend on the specific biometric trait, the matching algorithm, and the techniques used to fabricate the spoofed traits.

On the one hand, our results confirm the need of proper countermeasures against spoofing attacks, some of which have already been proposed in [5], [6]. On the other hand, they also show the need of novel methods to assess the robustness of multi-modal systems against spoofing attacks as the ones proposed by the authors in their preliminary works [8], [9], to overcome the limitations of the worst-case assumptions of [5], [6], [7].

Acknowledgment

This work was partly supported by the TABULA RASA project, 7th Framework Research Programme of the European Union (EU), grant agreement number: 257289; by the PRIN 2008 project "Biometric Guards - Electronic guards for protection and security of biometric systems" funded by the Italian Ministry of University and Scientific Research (MIUR); and by the Regione Autonoma della Sardegna, through the Regional Law n. 7 for Fundamental and Applied Research, in the context of the funded project Adaptive biometric systems: models, methods and algorithms, grant n. CP4 442. This work was also partly supported by a grant awarded to B. Biggio by Regione Autonoma della Sardegna, PO Sardegna FSE 2007-2013, L.R. 7/2007 "Promotion of the scientific research and technological innovation in Sardinia".

References

[1] X. Tan, Y. Li, J. Liu, L. Jiang, "Face Liveness Detection from A Single Image with Sparse Low Rank Bilinear Discriminative Model", *11th European Conference on Computer Vision*, pp. 504-517, 2010.

[2] T. Matsumoto, H. Matsumoto, K. Yamada and S. Hoshino, "Impact of Artificial "gummy" Fingers on Fingerprint Systems", *Op. Security and Counterfeit Deterrence Tech. IV, vol. 4677 of Proc. of SPIE*, pp. 275-289, 2002.

[3] X. He, Y. Lu and P. Shi, "A Fake Iris Detection Method Based on FFT and Quality Assessment", *Proc. Chinese Conf. on Pattern Recognition*, pp. 316-319, 2008.

[4] A. Ross, K. Nandakumar, A. K. Jain, "*Handbook of Multibiometrics*", Springer, 2006.

[5] R. N. Rodrigues, L.L. Ling, V. Govindaraju. "Robustness of multimodal biometric methods against spoof attacks". *J. of Visual Languages and Computing*, vol. 20, pp. 169-179, 2009.

[6] P. A. Johnson, B. Tan and S. Schuckers. "Multimodal Fusion Vulnerability to Non-Zero Effort (Spoof) Imposters". *Proc. IEEE Workshop on Information Forensics and Security*, pp. 1-5, 2010.

[7] R. N. Rodrigues, N. Kamat and V. Govindaraju. "Evaluation of Biometric Spoofing in a Multimodal System". *Proc. Fourth IEEE Int. Conf. Biometrics: Theory Applications and Systems*, pp. 1-5, 2010.

[8] Z. Akhtar, G. Fumera, G. L. Marcialis, and F. Roli. Robustness Evaluation of Biometric Systems under Spoof Attacks. *Proc. 16th Intl. Conf. on Image Analysis and Processing, 2011*, in press.

[9] Z. Akhtar, G. Fumera, G. L. Marcialis and F. Roli. "Robustness Analysis of Likelihood Ratio Score Fusion Rule for Multimodal Biometric Systems under Spoof Attacks". *Proc. 45th IEEE Int. Carnahan Conf. on Security Technology, 2011*, in press.

[10] B. Biggio, Z. Akthar, G. Fumera, G. L. Marcialis, and F. Roli. "Robustness of multi-modal biometric verification systems under realistic spoofing attacks". *Proc. Int. Joint Conf. on Biometric Systems, 2011*, in press.

[11] K. Nandakumar, Y. Chen, S. C. Dass and A. K. Jain. "Likelihood Ratio Based Biometric Score Fusion". *IEEE Tran. on Pattern Analysis and Machine Intelligence*, vol. 30, no. 2, pp.342-347, 2008.

[12] A. K. Jain, K. Nandakumar, A. Nagar. "Biometric Template Security". *EURASIP j. on Adv. in Sig. Proc.*, pp. 1-6, 2008.

[13] Y. Kim, J. Na, S. Yoon, and J. Yi. "Masked fake face detection using radiance measurements". *J. Opt. Soc. Am.* A 26, pp. 760-766, 2009.

[14] P. Coli, G.L. Marcialis, and F. Roli. "Vitality detection from fingerprint images: a critical survey". *IEEE/IAPR 2nd International Conference on Biometrics ICB 2007*, pp. 722-731, 2007.

[15] G.L. Marcialis, F. Roli. "Score-level fusion of fingerprint and face matchers under "stress" conditions". *In: Proc. Int. Conf. Image Analysis and Proc. (ICIAP)*, pp. 259-264, 2007.

[16] G.L. Marcialis, A. Lewicke, B. Tan, P. Coli, D. Grimberg, A. Congiu, A. Tidu, F. Roli, and S. Schuckers, "First International Fingerprint Liveness Detection Competition". *Proc. of 14th International Conference on Image Analysis and Processing ICIAP 2009*, pp. 9-12, 2009.

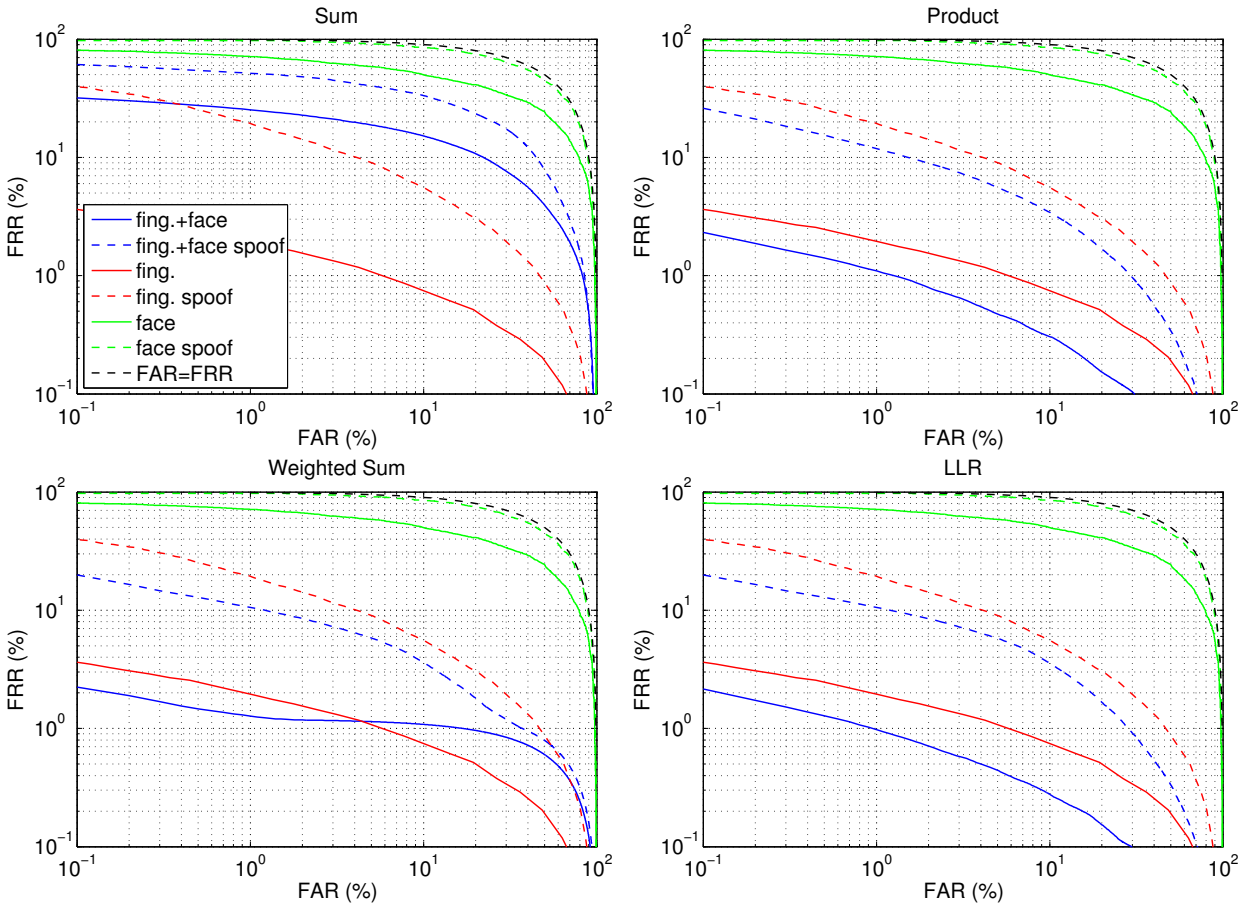

Fig. 4. Average DET curves obtained in our experiments on the testing set. Each plot refers to a different fusion rule, and contains DET curves of the systems under normal operation (solid curves) and under spoofing attacks (dashed curves). Green: mono-modal face system. Red: mono-modal fingerprint system. Blue: multi-modal face and fingerprint system (a spoofing attack against both traits is considered). Black: multi-modal system under a simulated worst-case spoofing attack against both traits.

[17] A. Godil, S. Ressler, P. Grother. "Face recognition using 3D facial shape and color map information: comparison and combination". *Biometric Technology for Human Identification, SPIE, vol. 5404*, pp. 351-361, 2005.

[18] Z. Zhang, D. Yi, Z. Lei, S. Z. Li. "Face Liveness Detection by Learning Multispectral Reflectance Distributions". *IEEE Int. Conf. on Automatic Face and Gesture Recognition and Workshops*, pp. 436-441, 2011.

[19] N. Poh and S. Bengio. "Using Chimeric Users to Construct Fusion Classifiers in Biometric Authentication Tasks: An Investigation". *Intl. Conf. on Acoustics, Speech, and Signal Processing*, pp. 1077-1080, 2006.

[20] http://www.nist.gov/itl/iad/ig/nbis.cfm

[21] http://www.cs.colostate.edu/evalfacerec/algorithms5.php