Anti-spoofing in action: joint operation with a verification system

Ivana Chingovska, André Anjos, Sébastien Marcel

{ivana.chingovska, andre.anjos, sebastien.marcel}@idiap.ch

Idiap Research Institute

Centre du Parc, Rue Marconi 19, CH-1920 Martigny

Abstract

Besides the recognition task, today's biometric systems need to cope with additional problem: spoofing attacks. Up to date, academic research considers spoofing as a binary classification problem: systems are trained to discriminate between real accesses and attacks. However, spoofing counter-measures are not designated to operate stand-alone, but as a part of a recognition system they will protect. In this paper, we study techniques for decision-level and score-level fusion to integrate a recognition and anti-spoofing systems, using an open-source framework that handles the ternary classification problem (clients, impostors and attacks) transparently. By doing so, we are able to report the impact of different spoofing counter-measures, fusion techniques and thresholding on the overall performance of the final recognition system. For a specific use-case covering face verification, experiments show to what extent simple fusion improves the trustworthiness of the system when exposed to spoofing attacks.

1. Introduction

From identity cards to personal computers' login information: biometrics has become a technology which supports many of our daily activities. Biometric recognition systems have been noticing constant improvement despite problems like illumination conditions, noisy equipment, wide range of viewpoints, aging of subjects etc. Nevertheless, the growing trend of mobile devices usage has popularized a new challenging problem: spoofing attacks. Spoofing is a non-zero effort attack and occurs when an invalid user tries to access a biometrically protected system by showing a copy of the biometric traits of a valid user. Since 2009, when the face verification systems of several laptops were successfully deceived by presenting fake facial images [10], every new commercial biometric verification system is being put to similar test by security enthusiasts.

The vulnerability to spoofing has been attested for many biometric modalities [18]. An attacker can spoof face

modality by showing a photo or a video of a valid user, fingerprint modality with gummy fingers, or iris modality with high-quality iris photographs. The biometric community constantly arrives with new solutions addressing more and more challenging spoofing attacks [27], [30], [7], [28], some of which are successful against a variety of attacks [15]. One piece missing in the puzzle seems to be how to put the anti-spoofing system where it belongs: working together with a recognition system.

By definition, a spoofing counter-measure is not designated to operate as a stand-alone system. Its purpose is to guard a biometric recognition system by performing additional checks whether the input comes from a live user or is a fake sample. An applicable biometric system is essentially a recognition system which may have increased robustness to spoofing due to being coupled with an anti-spoofing system. Yet, the publications on anti-spoofing tend to focus on the spoofing detection itself and omit to make the link to a recognition system. A joint operation with a recognition system is the setup where a spoofing counter-measure gets its meaning.

The previous observations throw light on two different problems. The first one is concerned with evaluation methodology. Integrating the counter-measure to a recognition system will affect the performance of the latter. In particular, it will probably reduce its vulnerability to spoofing attacks, but it may also disturb its recognition performance. Although evaluating an anti-spoofing system on its own is valid for comparing the effectiveness of different approaches against spoofing attacks, the final result that we are interested in is the performance of the recognition system. Once the danger of spoofing attacks is acknowledged, the recognition system has an additional class to discriminate besides impostors and valid users. Therefore, it needs to report three performance numbers instead of the two well known False Acceptance Rate (FAR) and False Rejection Rate (FRR). The second problem is related to the actual co-operation of the two systems. What is the best way to blend in an anti-spoofing system into a recognition system so that the final product is robust to spoofing, but does not suffer

from significantly reduced recognition accuracy?

The question of integration of a recognition (particularly in this paper, a verification) and anti-spoofing system for the face modality is the main subject of this work. We investigated four fusion strategies which accept or reject a verification attempt using the outputs both of a face verification and anti-spoofing system. For this purpose, we selected one baseline face verification system and three state-of-the-art counter-measures to face spoofing. We report the performance of the face verification algorithm before and after the fusion and we give a comparative analysis of the fusion rules using a suitable evaluation framework. As an additional contribution, we will provide an open-source implementation of the algorithms for easy reproduction of results.

The paper complies to the following structure: Section 2 delivers an overview of the fusion strategies used in biometrics and the sparse set of attempts for fusion of biometric recognition and anti-spoofing systems. Section 3 explains the fusion methodologies we investigated. The experimental results are presented in Section 4, after an introduction to the baseline systems involved in our experiments, as well as the database used for evaluation. Section 5 gives the conclusions of our work.

2. Related work

Fusion of multiple experts in biometrics is a well investigated field. It can remedy many limitations of the biometric systems that depend on a single trait, like sensitiveness to noisy data or failure-to-enroll problems. Even more, it can deliver better accuracy than each of the baseline systems alone [24]. In general, fusion strategies have been only applied to fuse multiple biometric experts.

The fusion of the biometric systems can happen at several points of the recognition pipeline. With regards to this, the fusion techniques are categorized as sensor-level, feature-level, score-level, rank-level and decision-level fusion. Having at disposal the outputs of several already developed systems, in this work we are going to focus on decision-level and score-level fusion.

The decision-level fusion strategies vary from simple ones like AND or OR rules, or majority voting, to more complex ones which rely on the probabilities of errors of the individual systems [24]. The score-level fusion can be performed using very simple score combinations, like sum or order statistics of scores, also referred to as *fixed fusion rules* [22]. These schemes require that the scores of the multiple experts are brought into a common domain by performing score normalization. More advanced schemes rely on Bayes decision theory: they compute the likelihoods that the set of multiple experts' scores belong to each of the classes and compare them using Likelihood Ratio (LLR). A nice theoretical foundation of these so-called *density-based fusion rules* [24], is given in [14]. Finally, many

systems use the scores as an input to a classifier who takes the final decision. Examples of these *learning-based fusion rules* are k-means or fuzzy clustering [6], Support Vector Machines (SVM) [4], [11], Linear Discriminant Analysis (LDA) [4], [23], [26], decision trees [4], [23], multi-layer perceptron [4], [26], Behavior Knowledge Space [22], [25] and many more.

The fusion rules can be applied to combine experts which work with the same biometric modality (like multi-sensor, multi-algorithm, multi-instance and multi-sample systems) or different biometric modalities (multi-modal) [24]. It may be intuitive to imagine that a multi-modal biometric system will be more robust to spoofing, because one needs to bring copies of more then one biometric trait to deceive the system. However, in many cases spoofing only one modality can be enough, as discussed in [21], [13], [1]. This is highly dependent on the used fusion algorithm, which has inspired fusion schemes specifically designed to increase the robustness to spoofing. Most notable are [21] which explores fuzzy logic, and [20] which extends the LLR scheme by introducing hidden variables denoting the probability of a spoofing attack. None of the mentioned work employs an algorithm specialized to detect spoofing attacks.

Prior work on fusion of biometric recognition and anti-spoofing system has not been as extensive. In fact, to the best of our knowledge, there are only two publications treating this topic and they are in the domain of fingerprint verification. In [17], an anti-spoofing mechanism is employed before the verification stage on a multibiometric system composed of three fingerprint and one face modality. If a spoofing attack is detected for one modality, the corresponding unimodal system does not contribute to the final score-level fusion of the multimodal system. On the other hand, [16] analyzes four different fusion methods. The first two sequentially employ a fingerprint verification and anti-spoofing system, or the other way around. The third method performs classification on the verification and anti-spoofing scores, while the fourth one creates a Bayesian Network which models a relationship between the verification and anti-spoofing scores. An important, but limiting assumption is that the enrollment dataset contains spoofing attacks. Another limitation is that the output of the third and fourth method are discrete, which prohibits graphical comparison with other methods.

3. Fusion strategies

A typical verification system on its own is trained to discriminate only between two classes: valid users as a positive and impostors as a negative class. If we plot the score distribution of a good verification system, it should be expected that the score distributions of the two classes are well separated. With the appearance of spoofing attacks, instead

of two, the system is now confronted with three classes: valid users, impostors and spoofing attacks. Luckily, the final system we are interested in needs to reject both impostors and spoofing attacks, and thus they can be considered as one enhanced negative class. However, if the spoofing attacks are of good quality, their score distribution may be close to, or even overlap the distribution of the valid users. As a consequence, the positive and the enhanced negative class are not that well separated any more.

To remedy this problem, in this work we investigated two approaches: decision-level and score-level fusion of verification system with an anti-spoofing system. The first approach regulates the decision taken by the verification system with an additional check performed by an anti-spoofing system. The second approach shifts the scores of the spoofing attacks in an attempt to create a margin between them and the score distribution of the valid users.

Unlike fusion of two biometric recognition experts, our task at hand requires fusing of two discordant systems: a verification and an anti-spoofing system are of different nature and have antagonistic criteria for taking a decision. As illustrated in Table 1, what is positive class for one system may be negative for the other and vice-versa. In particular, a verification system has to reject a zero-effort impostor, while an anti-spoofing system will probably recognize it as valid access. On the other hand, a verification system may consider a spoofing attack sample as a positive class, while the anti-spoofing system is trained to reject it.

	valid users	impostors	spoofs
Verification	+	-	+
Anti-spoofing	+	+	-
Final system	+	-	-

Table 1. Criteria for positive and negative class of a typical verification and anti-spoofing system and the final system of interest

Table 1 may give hints about the fusion schemes that we can employ. In the case of the decision-level fusion approach, the positive class needs to be accepted by both the verification and anti-spoofing system, while for the negative class a rejection from one of the systems is enough. Thus, we fuse the decisions of the two systems using AND fusion rule [24].

In the case of score-level fusion, if we note that each of the separate systems gives high scores to the class it considers as positive and low scores to the class it considers as negative, then many schemes typically used in biometrics may not be straight-forwardly applicable in this case. To illustrate the reason, one could consider the classical Bayesian approaches, like the product or the sum rule of posterior probability of the class given the evidence [14]. Let us imagine the response of the two systems under a spoofing attack. The posterior probability that the spoofing


Figure 1. Fusion of verification and anti-spoofing algorithm

attack is a positive sample may be very high for the verification classifier, but very low for the spoofing detector. On the contrary, it can be expected that the posterior probability of a zero-effort impostor will be low for the verification classifier, and high for the spoofing detector. As a result, in many cases the final output may be ambiguous and the final decision wrong.

We investigated three score-level fusion mechanisms and how they affect the performance of the verification system. Given the verification score s_v and the anti-spoofing score s_s of an input sample, we generate a fused score s_f which combines them together in an appropriate way.

The first explored fusion option belongs to the class of fixed fusion rules and performs simple sum of the scores. As an alternative, we consider the scores of the two separate systems as points in a 2D space, where the first dimension is s_v and the second dimension is s_s. After an analysis of the distribution of the points belonging to the separate three classes in the 2D space, we can train a suitable learning-based fusion rule to distinguish between score pairs of the ultimate positive class (valid users) and the two classes of negatives (impostors and spoofing attacks). We selected Logistic Regression (LR) [8] as the first learning based fusion scheme. Given the 2D data of the training set as input variables, we fit a logistic hypothesis function given in Eq. 1 by estimating its parameters Θ. The parameters Θ are in a linear relation with the input variables i.e. $\mathbf{x} = \begin{pmatrix} s_v & s_s \end{pmatrix}^\top$ in Eq. 1.

$$h_\Theta(\mathbf{x}) = \frac{1}{1 + e^{-\Theta^\top \mathbf{x}}} \qquad (1)$$

Besides LR, we performed experiments with Polynomial Logistic Regression (PLR), where the parameters Θ are in a polynomial relation with the input variables, i.e. $\mathbf{x} = \begin{pmatrix} s_v & s_s & s_v s_s & s_v^2 & s_s^2 \end{pmatrix}^\top$ in Eq. 1. The motivation lies on the assumption of non-linear separability between the positive and the negative class.

The workflow of a fused system is illustrated in Figure 1. Depending on the fusion strategy, the fusion module either computes two decision thresholds directly on the separate systems' output scores s_v and s_s (decision-level fusion), or computes a single decision threshold on the combined score

s_f (score-level fusion).

4. Experimental results

The selected use-case scenario for empirical evaluation of the discussed fusion mechanisms is in the domain of face verification. Before proceeding with the results, in Section 4.1 we give an introduction to the systems that we used for our experiments: spoofing database, baseline face verification and anti-spoofing system. As a first experiment, in Section 4.2 we report the performance of the face verification algorithm alone, and justify the need of spoofing counter-measure. Then, in Section 4.3 we proceed with comparing the fusion algorithms with regards to the anti-spoofing algorithm used. All the experiments are coded using the free signal-processing and machine-learning toolbox Bob[1] [2]. The source code of our work is available as free software to provide easy reproduction of results[2]

4.1. Spoofing database and baseline systems

Extensive experiments illustrating the outcomes of fusion between face verification and anti-spoofing systems are performed in a realistic scenario using a state-of-the-art face verification algorithm as a baseline and several face anti-spoofing algorithms. An important remark here is that, in order to train a baseline face verification system, we need a spoofing database which contains enrollment data which will be used to create the models of the clients. Then, to obtain the face verification scores for the real accesses, we match each of the samples against each of the models of the enrolled clients in an exhaustive manner. Depending on whether the model and the sample belong to the same identity or not, we get score sets for valid users and impostors. This set of matches compose the so-called *licit* protocol, which resembles the typical protocol used for evaluation of verification systems. To obtain the face verification scores for the spoofing attacks, we match each of the spoofing samples to the models of the corresponding client. This set of matches compose the so-called *spoof* protocol. While the licit protocol will be used to assess the verification performance of the algorithm, the spoof protocol will be used to evaluate its performance when spoofing attacks are present. Note that the positives of both licit and spoof protocol are the same. Only the licit protocol is involved in the development stage, i.e. in determining the decision threshold.

A face spoofing database which delivers enrollment data, and defines precise training, development and test set for unbiased evaluation of algorithms, is Replay-Attack [7]. It contains three types of face spoofing attacks (printed photos, digital photos displayed on a screen and video attacks) for a total of 50 identities. Additional advantage is the

data in the form of video sequences, which allows for experimentation with liveness or motion based anti-spoofing methods requiring temporal data to detect spoofing attacks.

The baseline face verification system is based on [5]. From the input face images, it extracts Discrete Cosine Transform (DCT) features and builds Gaussian Mixture Model (GMM) as a Universal Background Model. To train such a system and obtain the face verification scores, we used the implementation available in the open-source face verification framework from [12].

Regarding the face anti-spoofing methods, we experimented with three different algorithms. Two of them are motivated by the assumption that the face spoofing attacks and the real accesses have differences in their texture patterns, due to the possible artifacts and image quality degradation that may appear in the printing or recapturing procedure. In particular, the first algorithm, given in [7], uses uniform Local Binary Patterns (LBP) [19] to capture the texture properties and builds the feature vector based on the LBP histograms. The second algorithm [9] uses an LBP variant which extends the operator into temporal dimension, thus capturing dynamic texture properties in three orthogonal planes (LBP-TOP) [29]. The third algorithm [3] analyzes the motion patterns present in the scene. It is justified by the observation that the motion patterns on the face and in the background of an input sample will be correlated in the case of spoofing attack and de-correlated in the case of a real access.

4.2. Performance of the face verification system

As stated in Section 1, the presence of spoofing attacks implies that we report three error rates when evaluating a face verification system. Among the many possibilities, we adapt the proposition of [13]. Besides FAR (the ratio of wrongly accepted impostors), FRR (the ratio of wrongly accepted impostors), and HTER (Half Total Error Rate) as an average between FAR and FRR, [13] introduces Spoof False Acceptance Rate (SFAR), which represents the ratio of successful spoofing attacks among all the available spoofing attacks.

After training the GMM-based baseline face verification system using the training set and determining a decision threshold at the point of Equal Error Rate (EER) for the development set of Replay-Attack, in Figure 2(a) we plot the verification scores of the system for the three groups of samples (valid users, impostors and spoofing attacks) for the test set, as well as the decision threshold taken at EER on the development set. The full green line represents the probability of a spoofing attack deceiving the system depending on the chosen threshold.

Despite the great performance of the baseline in the face verification task (**FAR=0.05%; FRR=0.24%; HTER=0.14%**), the system is severely deceived when ex-

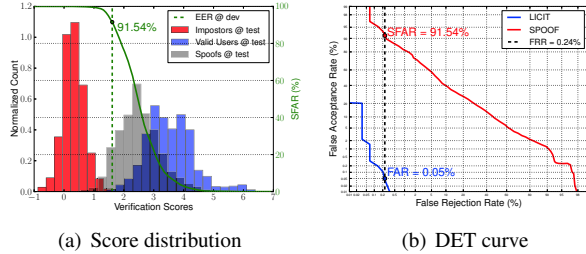(a) Score distribution   (b) DET curve

Figure 2. GMM-based face verification system: performance figures

posed to spoofing attacks. The spoofing attacks of Replay-Attack are closely resembling the real accesses, at least from the perspective of the analyzed baseline. Indeed, the verification scores of Replay-Attack are in the range of the real accesses, resulting in **91.54%** of the spoofing attacks successfully intruding into the system. This demonstrates the necessity of an anti-spoofing system to protect the face verification baseline.

Additionally, we plot DET curve of the verification performance of the baseline system using the licit protocol in Figure 2(b). We overlay that plot with the DET curve of the spoof protocol. As expected, due to the low separability of the valid users and spoofing attacks, the overlaid DET curve is much above DET for licit protocol. Since the positives of the licit and the spoof protocol overlap as explained in Section 4.1, FRR will be the same for the two protocols no matter what the selected threshold is. Figure 2(b) shows FAR and SFAR that correspond to FRR=0.24% when the decision threshold is taken at the EER of the development set. We hope that securing the baseline system by fusion with anti-spoofing system will reduce the space between the two DET curves with as little effect on the licit protocol curve as possible.

4.3. Performance of the fused systems

With our experiments, we wanted to observe how fusion with an anti-spoofing system affects the performance of the baseline face verification system. Table 2 gives HTER on the licit protocol and SFAR on the spoof protocol for a fused system using the four fusion strategies. The AND fusion rule requires estimation of two decision thresholds: one for the verification and one for the anti-spoofing system. The thresholds are taken at EER on the development set. For the score-level fusion strategies, each sample has a single fused score and hence a single threshold needs to be estimated. It is taken at EER on the development set using the licit protocol. The results are reported on the test set. Column-wise, one can compare how the choice of anti-spoofing algorithm affects the system's performance under

fusion with particular fusion mechanism. Row-wise, comparison of different fusion mechanisms for systems using identical counter-measure is possible.

In comparison to SFAR = 90.54% of the baseline face verification system, Table 2 shows how the system benefits from fusion: in all the fused systems SFAR is significantly reduced. It can go down to as much as only 3.62% when the fusion is done with LBP-TOP and SUM fusion rule. Comparing the anti-spoofing algorithms with regards to the achieved SFAR, LBP-TOP boosts the most the system's robustness to spoofing. As for the fusion mechanisms, SUM fusion outperforms the other fusion rules in rejecting spoofing attacks.

However, the fused systems can not be evaluated solely based on the reduction of SFAR. Decreasing SFAR almost always comes with the cost of HTER significantly higher than the baseline. For example, SUM fusion records the greatest deterioration of HTER, although its SFAR is lower then the other fusion rules. On the other hand, PLR fusion is not as effective in rejecting spoofing attacks, but the verification performance of the system is almost unaffected. Lacking a heuristic to determine which system performs the best based on the values of HTER and SFAR, one needs to choose the system that provides the required trade-off between these two values.

While the trade-off paradigm is valid when comparing the fusion rules, choosing the best anti-spoofing algorithm is not as difficult. LBP-TOP is advantageous with regards to both HTER and SFAR almost in all the cases.

For a visual comparison of the four types of fusion algorithms with LBP-TOP counter measure, in Figure 3 we plot their decision boundaries. Each of the plots shows the distribution of the score points for the samples in the test set: the x coordinate of a point is its face verification score, while the y coordinate is its anti-spoofing score. The trade-off between SFAR and HTER is once again observable in these plots.

Choosing the setup of a PLR fusion system composed of GMM-based face verification baseline and LBP-TOP based anti-spoofing method, in Figure 4 we graphically present the performance changes of the baseline face verification system introduced by fusion. Comparing Figure 4(a) with Figure 2(a), we can observe how fusion shifts the scores of the spoofing attacks towards the impostors, creating larger margin between them and the valid users. Figure 4(b) shows how the DET lines for the licit and spoof protocol appear to be closer to each other compared to the baseline in Figure 2(b). This results in a more balanced trade-off between the verification performance and the robustness to spoofing for the final system, regardless of the chosen threshold.

	AND		SUM		LR		PLR	
	HTER	SFAR	HTER	SFAR	HTER	SFAR	HTER	SFAR
LBP	11.04	8.93	13.74	6.77	5.17	15.05	3.08	27.1
LBP-TOP	6.46	4.00	7.79	3.62	2.75	9.02	1.48	13.67
MOTION	4.74	11.78	10.31	9.1	2.26	35.34	1.75	39.61

Table 2. Performance table of fused GMM-based face verification and anti-spoofing systems (in %)

(a) AND fusion   (b) SUM fusion   (c) LR fusion   (d) PLR fusion

▲ ▲ impostors   ■ ■ spoofing attacks   ● ● real accesses   ▬ decision boundary
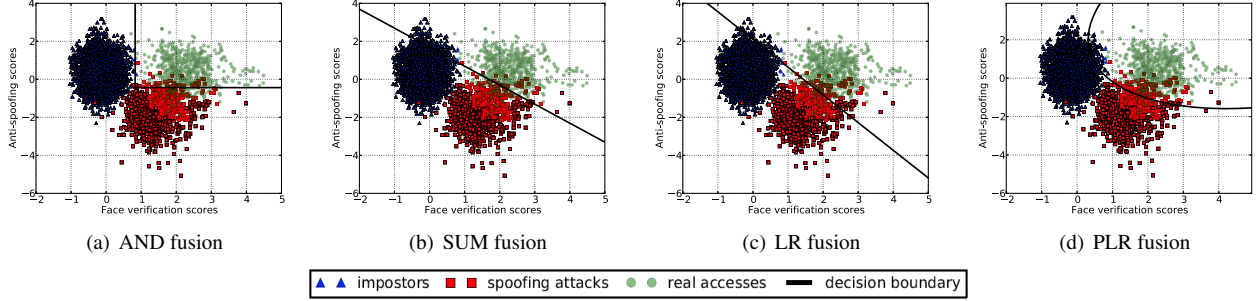
Figure 3. Decision boundaries for different fusion rules on the scatter plot of the scores of GMM-based face verification and LBP-TOP anti-spoofing algorithm


(a) Score distribution   (b) DET curve

Figure 4. PLR fused system between GMM-based face verification and LBP-TOP based anti-spoofing system: performance figures

5. Conclusions

Although the research in anti-spoofing records more and more notable achievements in detecting spoofing attacks, few of the spoofing detection schemes have been put in a practical setup where they have to cooperate with a recognition system. An experiment of this kind is very important because an anti-spoofing algorithm is not expected to work on its own and the integration may significantly influence the performance of the recognition system. This paper presents an approach to merge an anti-spoofing system to a face verification system by decision-level and score-level fusion. Four fusion rules have been tested: AND, SUM, LR and PLR. The performance of the fused face verification system is reported by the standard HTER, as well as SFAR as a ratio of spoofing attacks passing the system. The exper-

iments are conducted using a state-of-the-art GMM-based face verification system and three different anti-spoofing methods.

The typical behavior is that there is always a trade-off between HTER and SFAR of the system. Yet, fusion manages to reduce the ratio of successful spoofing attacks from 91.54% for the particular baseline face verification system, to less then 10% in some cases. The most successful in rejecting spoofing attacks is the system created with SUM fusion rule, but at the worst in terms of verification performance. The PLR fusion scheme almost perfectly preserves the verification capabilities of the baseline system, and still achieves good robustness to spoofing. It is up to the user to decide what trade-off of HTER vs. SFAR is convenient for the given application.

Although the particular use-case presented in this paper was in face verification, the analyzed fusion mechanisms are general and can be deployed for other modalities. The studied fusion schemes do not exhaust all the fusion possibilities. Besides the vast amount of decision-level and score-level fusion mechanisms typically used in biometrics, many of which are classifier-based, algorithms specifically designed for fusion of two systems of different nature can be attempted as well. Furthermore rank-level, or even feature-level fusion may be inspected. Regardless of the fusion technique, we believe that the anti-spoofing algorithms should not be evaluated exclusively in an isolated fashion. Rather the recognition systems which are protected by an anti-spoofing algorithm should be evaluated both for recognition performance and robustness to spoofing.

References

[1] Z. Akhtar, G. Fumera, G.-L. Marcialis, and F. Roli. Evaluation of serial and parallel multibiometric systems under spoofing attacks. In *5th IEEE International Conference on Biometrics: Theory, Applications and Systems*, 2012. 2

[2] A. Anjos et al. Bob: a free signal processing and machine learning toolbox for researchers. In *20th ACM Conference on Multimedia Systems (ACMMM)*. ACM Press, Oct. 2012. 4

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

[4] S. Ben-Yacoub, Y. Abdeljaoued, and E. Mayoraz. Fusion of face and speech data for person identity verification. *IEEE Transactions on Neural Networks*, 10. 2

[5] F. Cardinaux, C. Sanderson, and S. Marcel. Comparison of MLP and GMM classifiers for face verification on XM2VTS. In *Proceedings of the 4th International Conference on AVBPA*, 2003. 4

[6] V. Chatzis, A. G. Borş, and I. Pitas. Multimodal decision-level fusion for person authentication. *IEEE Transactions on Systems, Man and Cybernatics - Part A: Systems and Humans*, 29. 2

[7] I. Chingovska, A. Anjos, and S. Marcel. On the effectiveness of local binary patterns in face anti-spoofing. In *Proceedings of the 11th International Conference of the Biometrics Special Interes Group*, 2012. 1, 4

[8] W. H. David and L. Stanley. *Applied Logistic Regression*. Wiley-Interscience Publication, 2000. 3

[9] T. de Freitas Pereira, A. Anjos, J. M. De Martino, and S. Marcel. LBP-TOP based countermeasure against face spoofing attacks. In *International Workshop on Computer Vision With Local Binary Pattern Variants - ACCV*, page 12, 2012. 4

[10] N. M. Duc and B. Q. Minh. Your face is not your password. Black Hat Conference, 2009. 1

[11] J. Fierrez-Aguilar, J. Ortega-Garcia, D. Garcia-Romero, and J. Gonzalez-Rodriguez. A comparative evaluation of fusion strategies for multimodal biometric verification. In *Fourth International Conference on Audio- and Video-based Biometric Person Authentication*, pages 830–837, 2003. 2

[12] M. Günther, R. Wallace, and S. Marcel. An open source framework for standardized comparisons of face recognition algorithms. In *Computer Vision - ECCV 2012. Workshops and Demonstrations*, volume 7585, pages 547–556, 2012. 4

[13] P. A. Johnson, B. Tan, and S. Schuckers. Multimodal fusion vulnerability to non-zero (spoof) imposters. In *IEEE International Workshop on Information Forensics and Security*, 2010. 2, 4

[14] J. Kittler, M. Hatef, R. P. W. Duin, and J. Matas. On combining classifiers. *IEEE Transactions on Pattern Analysis and Machine Inteligence*, 20:226–239, 1998. 2, 3

[15] J. Komulainen et al. Complementary countermeasures for detecting scenic face spoofing attacks. In *International Conference on Biometrics (ICB'13)*, 2013. 1

[16] E. Marasco, Y. Ding, and A. Ross. Combining match scores with liveness values in a fingerprint verification system. In *5th IEEE International Conference on Biometrics: Theory, Applications and Systems*, 2012. 2

[17] E. Marasco, P. Johnson, C. Sansone, and S. Schuckers. Increase the security of multibiometric systems by incorporating a spoofing detection algorithm in the fusion mechanism. In *Proceedings of the 10th international conference on Multiple classifier systems*, pages 309–318, 2011. 2

[18] K. A. Nixon, V. Aimale, and R. K. Rowe. *Handbook of Biometrics*, chapter Spoof Detection Schemes. Springer-Verlag, 2008. 1

[19] T. Ojala, M. Pietikäinen, and T. Mäenpää. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(7):971–987, 2002. 4

[20] R. Rodrigues, N. Kamat, and V. Govindaraju. Evaluation of biometric spoofing in a multimodal system. In *Biometrics: Theory Applications and Systems (BTAS), 2010 Fourth IEEE International Conference on*, 2010. 2

[21] R. N. Rodrigues, L. L. Ling, and V. Govindaraju. Robustness of multimodal biometric fusion methods against spoofing attacks. *Journal of Visual Languages and Computing*, 20(3):169–179, 2009. 2

[22] F. Roli, J. Kittler, G. Fumera, and D. Muntoni. An experimental comparison of classifier fusion rules for multimodal personal identity verification systems. In *Proceedings of the Third International Workshop on Multiple Classifier Systems*, pages 325–336, 2002. 2

[23] A. Ross and A. Jain. Information fusion in biometrics. *Pattern Recognition Letters*, 24, 2003. 2

[24] A. Ross, K. Nandakumar, and A. K. Jain. *Handbook of Biometrics*, chapter Introduction to multibiometrics. Springer-Verlag, 2008. 2, 3

[25] M. Vatsa, R. Singh, A. Noore, and A. Ross. On the dynamic selection of biometric fusion algorithms. *IEEE Transactions on Information Forensics and Security*, 5(3):470–479, 2010. 2

[26] Y. Wang, T. Tan, and A. Jain. Combining face and iris biometrics for identity verification. In *Fourth International Conference on Audio- and Video-based Biometric Person Authentication*, 2003. 2

[27] D. Yambay et al. LivDet 2011 - fingerprint liveness detection competition 2011. In *Biometrics (ICB), 2012 5th IAPR International Conference on*, pages 208–215, 2012. 1

[28] H. Zhang et al. Learning hierarchical visual codebook for iris liveness detection. In *International Joint Conference on Biometrics*, 2011. 1

[29] G. Zhao and M. Pietikäinen. Dynamic texture recognition using local binary patterns with an application to facial expressions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(6):915–928, 2007. 4

[30] Z. Zhiwei et al. A face antispoofing database with diverse attacks. In *Proceedings of the 5th IAPR International Conference on Biometrics (ICB'12), New Delhi, India*, 2012. 1