

Multimodal Biometric Fusion: A Study on Vulnerabilities to Indirect Attacks

Marta Gomez-Barrero, Javier Galbally,
Julian Fierrez, and Javier Ortega-Garcia

Biometric Recognition Group–ATVS, EPS, Universidad Autonoma de Madrid,
C/ Francisco Tomas y Valiente 11, 28049 Madrid, Spain
{marta.barrero,javier.galbally,julian.fierrez,javier.ortega}@uam.es

Abstract. Fusion of several biometric traits has traditionally been regarded as more secure than unimodal recognition systems. However, recent research works have proven that this is not always the case. In the present article we analyse the performance and robustness of several fusion schemes to indirect attacks. Experiments are carried out on a multimodal system based on face and iris, a user-friendly trait combination, over the publicly available multimodal Biosecure DB. The tested system proves to have a high vulnerability to the attack regardless of the fusion rule considered. However, the experiments prove that not necessarily the best fusion rule in terms of performance is the most robust to the type of attack considered.

Keywords: Security, vulnerabilities, multimodality, iris recognition, face recognition, fusion schemes.

1 Introduction

Being able to automatically recognise people is of the utmost importance for many applications, such as regulating international border crossings or performing financial transactions on-line. Traditional security technologies required the use of PINs or tokens. Biometrics proposes a change of paradigm, from “what you know” or “what you have” to “who you are”: forget about passwords, you are your own key [1].

However, as any other security technology, biometrics are exposed to external attacks which could compromise their integrity [2]. It is therefore essential to understand the threats to which they are subjected and to analyse their vulnerabilities in order to prevent possible attacks and increase their benefits for the users. External attacks to biometric systems are commonly divided into: *direct attacks* (also known as *spoofing attacks*), carried out against the sensor, and *indirect attacks*, directed to some of the inner modules of the system. In the last recent years important research efforts have been conducted to study the vulnerabilities of biometric systems to both direct and indirect attacks [3–5].

This new concern which has arisen in the biometric community regarding the security of biometric systems has led to the appearance of several international

projects, like the European TABULA RASA [6] and BEAT [7], which base their research on the security through transparency principle: in order to make biometric systems more secure and reliable, their vulnerabilities need to be analysed and useful countermeasures need to be developed.

In this scenario, biometric multimodality has been regarded as an effective way of increasing the robustness of biometric-based security systems against external attacks. Combining the information offered by several traits would force an eventual intruder to successfully break several unimodal modules instead of just one. However, it has already been proven that this is not necessarily true for the case of spoofing attacks [8–10].

But are all fusion schemes equally robust to indirect attacks? If not, are the system performance and the robustness somehow correlated? In the present work we try to answer those questions using several score-level fusion schemes and a multimodal indirect attack already proven to be very successful in [11].

The paper is structured as follows: the attacking algorithm is summarized in Sect. 2. The system attacked, with the different fusion rules considered, is presented in Sect. 3, while the experimental protocol followed and the performance evaluation of the system are described in Sect. 4. The results obtained are shown in Sect. 5. Finally conclusions are drawn in Sect. 6.

2 Hill-Climbing Attack to Multimodal Recognition Systems

In order to attack the multimodal verification system using the different fusion schemes considered, the algorithm detailed in [11] will be used, which may be summarized as follows. Consider the problem of finding a $(K + L)$ -dimensional vector \mathbf{x} of real (size K) and binary (size L) values which, compared to an unknown template \mathcal{C} (in our case related to a specific client), produces a similarity score higher than a certain threshold δ , according to some matching function J , i.e., $J(\mathcal{C}, \mathbf{x}) > \delta$.

The problem stated above may be solved by dividing the vector \mathbf{x} into its real-valued (\mathbf{x}_{real}) and binary parts (\mathbf{x}_{bin}) and alternately optimizing each of them. In order to optimize each of the parts, two different sub-algorithms will be used: *i*) a hill-climbing based on the uphill simplex to attack the real-valued segment; and *ii*) a hill-climbing attack based on a genetic algorithm to break the binary segment. Thus, the steps followed by the multimodal attack are:

1. Generate a synthetic template (\mathbf{x}) randomly initializing the real-valued (\mathbf{x}_{real}) and binary (\mathbf{x}_{bin}) segments, of lengths K and L , respectively. Then compute the similarity score $s = J(\mathcal{C}, \mathbf{x})$, which will be iteratively maximised.
2. Leaving one of the segments unaltered, optimize the other segment of the template using the appropriate sub-algorithm until one of the stopping criteria of the sub-algorithm is fulfilled.
3. Change the optimization target to the segment which was previously left unaltered and go back to step 2.

The algorithm stops when: *i*) the verification threshold is reached (i.e., access to the system is granted), or *ii*) the total number of iterations exceeds a previously fixed value (i.e., the attack has failed).

It should be noted that the number of executions of each sub-algorithm is not fixed, and may vary depending on the user account at hand. That number can even be zero for one of the sub-algorithms, meaning that optimizing the other part of the template is enough to break the account.

For further details on the multimodal attack and on each of the two sub-algorithms, the reader is referred to [11].

Notation. Since the multimodal attack will be tested against a face- and iris-based multimodal system, we will henceforth denote the number of times the real-valued hill-climbing is executed as N_{face} , and the number of times that the binary-valued hill-climbing is executed as N_{iris} . Similarly, the real-valued segment of the template \mathbf{x} will be denoted as \mathbf{x}_{face} , and the binary part as \mathbf{x}_{iris} .

3 Multimodal Verification System

The multimodal verification system evaluated in this work is the fusion of two unimodal systems, namely: *i*) the iris recognition system developed by L. Masek¹ [12], which is widely used in related publications; and *ii*) an Eigenface-based face verification system, used, for instance, to present initial face verification results for the Face Recognition Grand Challenge [13].

Given an input vector \mathbf{x} , the multimodal system performs the following tasks in order to obtain the final score, s :

1. Compute the similarity scores obtained by the face (s_{face}) and iris (s_{iris}) traits, as given by the unimodal matchers.
2. Normalize the scores s_k , with $k = \{\text{face}, \text{iris}\}$, using hyperbolic tangent estimators (its robustness and high efficiency are proven in [14]). This way, the normalised scores s'_k lie in the interval $[0, 1]$.
3. Finally, both normalised scores are fused. Several fusion schemes have been considered [15, 16]:

$$\begin{array}{ll} \text{Sum rule : } s = s'_{\text{face}} + s'_{\text{iris}} & \text{Product rule : } s = s'_{\text{face}} \times s'_{\text{iris}} \\ \text{Max rule : } s = \max\{s'_{\text{face}}, s'_{\text{iris}}\} & \text{Min rule : } s = \min\{s'_{\text{face}}, s'_{\text{iris}}\} \end{array}$$

4 Database and Experimental Protocol

The experiments are carried out on the face and iris subcorpora included in the Desktop Dataset of the Multimodal Biosecure Database [17], which comprises voice, fingerprints, face, iris, signature and hand of 210 users, captured in two time-spaced acquisition sessions. This database was acquired thanks to the joint

¹ www.csse.uwa.edu.au/~pk/studentprojects/libor/sourcecode.html




Fig. 1. Typical samples of the face and iris images available in the Desktop Dataset of the multimodal BioSecure database

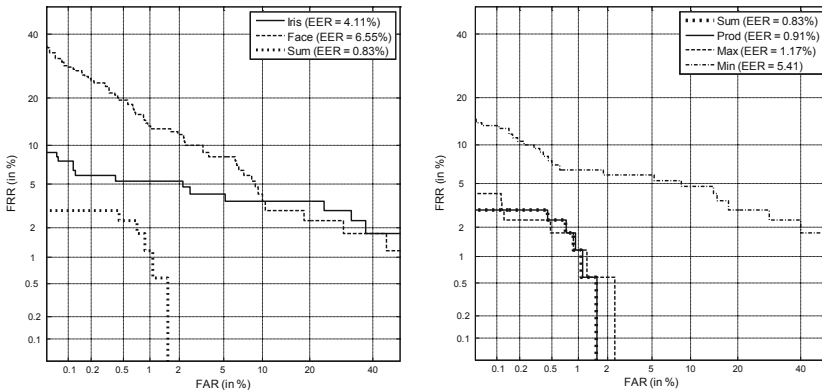


Fig. 2. DET curves for the unimodal systems and the fusion rule with the best performance (left) and for all the fusion rules considered (right), with their corresponding EER

effort of 11 European institutions and has become one of the standard benchmarks for biometric performance and security evaluations. It is publicly available through the BioSecure Association².

The face subset used in this work includes four frontal images (two per session) with an homogeneous grey background, and captured with a reflex digital camera without flash ($210 \times 4 = 840$ face samples), while the iris subset includes four grey-scale images (two per session as well) per eye, all captured with the Iris Access EOU3000 sensor from LG. In the experiments only the right eye of each user has been considered, leading this way as in the face case to $210 \times 4 = 840$ iris samples. Typical samples may be seen in Fig. 1.

4.1 Performance Evaluation

The database is divided into: *i*) a training set comprising the first three samples of 170 clients, used as enrolment templates for each sub-system; and *ii*) an evaluation set formed by the fourth image of the previous 170 users (used to

² <http://biosecure.it-sudparis.eu/AB>

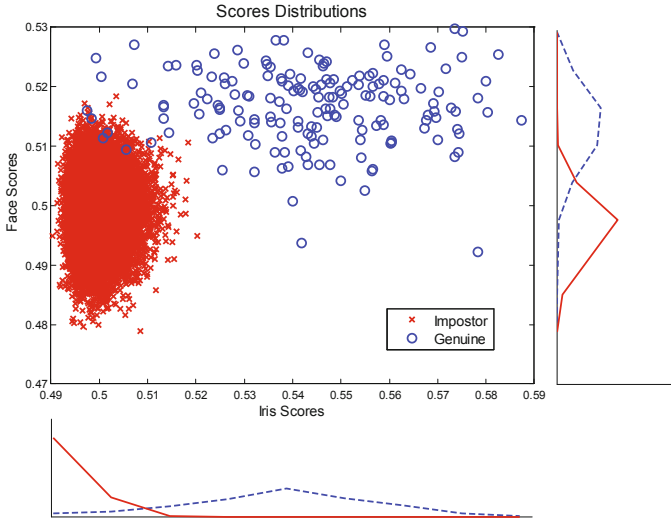


Fig. 3. Genuine and impostor distributions for the face (y axis) and iris (x axis) recognition systems

compute the genuine scores), and all the 4 images of the remaining 40 users (used to compute the impostor scores). The final score given by the multimodal system is the average of the scores obtained after matching the input template \mathbf{x} to the three face and iris templates of the client model \mathcal{C} .

The attacking algorithm is evaluated at three operating points with FAR = 0.1%, FAR = 0.05%, and FAR = 0.01%, which correspond to a low, medium, and high security application according to [18].

As described in Sect. 3, several fusion rules are considered in the present study. The verification performance of the unimodal and multimodal combinations considered are shown in Fig. 2, where the Detection Error Tradeoff (DET) curves are depicted. As may be observed, the best performance is achieved for the sum rule (EER = 0.83%), while the worst one is shown for the min rule (EER = 5.41%).

In Fig. 3, the genuine and impostor distributions are shown.

4.2 Experimental Protocol for the Attacks

The performance of the attack will be evaluated in terms of: *i*) Success Rate (SR) which is the expected probability of breaking a given account, indicating how dangerous the attack is (the higher the SR, the bigger the threat); and *ii*) Efficiency (Eff) defined as the inverse of the average number of matchings needed to break an account, thus giving an estimation of how easy it is for the attack to break into the system in terms of speed (the higher the Eff, the faster the attack). The SR is computed as the ratio between the number of broken accounts (A_B) and the total number of accounts attacked ($A_T = 170$): $SR = A_B/A_T$, and the

Table 1. Eff and SR for the different fusion rules considered

FAR	Sum		Prod		Max		Min	
	SR	Eff ($\times 10^{-4}$)	SR	Eff ($\times 10^{-4}$)	SR	Eff ($\times 10^{-4}$)	SR	Eff ($\times 10^{-4}$)
0.10%	100%	1.9372	100%	1.9144	100%	1.3231	100%	2.3134
0.05%	100%	1.8218	100%	1.7863	100%	1.2060	100%	2.0602
0.01%	100%	1.3702	100%	1.3616	100%	1.0220	100%	1.7657

Table 2. Number of user accounts broken after attacking each part of the template a fixed number of times specified by N_{face} and N_{iris} (see Sect. 2)

FAR	Sum ($N_{face} + N_{iris}$)						Prod ($N_{face} + N_{iris}$)					
	1+0	1+1	2+1	2+2	3+2	3+3	1+0	1+1	2+1	2+2	3+2	3+3
0.10%	0	153	9	7	0	1	0	161	5	4	0	0
0.05%	0	155	8	7	0	0	0	158	6	6	0	0
0.01%	0	117	27	21	2	3	0	118	27	19	3	3

FAR	Max ($N_{face} + N_{iris}$)						Min ($N_{face} + N_{iris}$)					
	1+0	1+1	2+1	2+2	3+2	3+3	1+0	1+1	2+1	2+2	3+2	3+3
0.10%	5	118	14	30	1	2	0	127	19	4	13	0
0.05%	2	102	15	36	4	9	0	101	37	7	20	0
0.01%	0	90	5	54	8	10	0	58	53	8	38	0

Eff is computed as $Eff = 1 / \left(\sum_{i=1}^{A_B} n_i / A_B \right)$, where n_i is the number of matchings needed to bypass each of the broken accounts.

5 Results

The experiments have two different goals, namely: *i*) analyse the robustness against indirect attacks of different fusion rules, and *ii*) study to what extent the vulnerabilities of a multimodal recognition system based on face and iris are correlated to the verification performance.

5.1 Vulnerabilities Evaluation

In Table 1, the performance of the attack in terms of the SR and the Eff is shown. As may be observed, the SR is 100% in all cases: all accounts are broken, regardless of the fusion scheme considered. However, not all the fusion schemes are equally robust in terms of speed: the Eff for the min rule is the highest one, being therefore the least robust fusion scheme. On the other hand, while the Eff for the sum and product rules is very similar, for the max rule it is considerably lower. Therefore, for applications where the robustness to this kind of attacks is more important than having an optimal performance (EER rises from 0.83% with the sum rule, to 1.17% with the max rule), the max rule should be considered.

For all the user accounts attacked, each sub-algorithm was executed between 0 and 3 times. Therefore, there are six possible cases regarding the number of those executions ($N_{\text{face}} + N_{\text{iris}}$). In the particular case when $N_{\text{iris}} = 0$, the account was broken after the first execution of the real-valued hill-climbing, therefore not needing to attack the binary part. The number of accounts that fall into each category is shown in Table 2. As may be observed, most accounts are broken after optimizing each part of the template only once.

In Sect. 4.2, Eff was defined as the inverse of the average number of comparisons needed to break an account. Therefore, the lower the Eff, the higher the number of comparisons needed. As could be expected, the lower the FAR at the operating point tested, the higher the number of users for which more executions of each sub-algorithm were needed.

However, when we compare the results shown in Table 2, we observe two different behaviours:

- For the sum, product and max rules, as expected, the lower the Eff, the higher the number of users for which two or even three executions of each sub-algorithm were needed.
- For the min rule, which presented the highest Eff for the attack (see Table 1), the number of users requiring three executions of the real-valued sub-algorithm is the highest. This means that the genetic sub-algorithm saturates quickly, and therefore the general attacking scheme starts attacking the face part of the template: as stated in [11], the genetic sub-algorithm needs considerably more comparisons than the hill-climbing based on the uphill simplex, leading this quick change to a higher Eff.

6 Conclusions

In the present article we have analysed the robustness of different multimodal score-level fusion rules (sum, product, max and min) to indirect attacks. We have then explored to what extent there is a correlation between the vulnerability level and the performance of the multimodal system. A multimodal system based on face and iris, a trait combination commonly regarded as user-friendly, working on a publicly available multimodal database, was used in the experiments.

The experiments showed that the multimodal attack achieves a Success Rate of 100% in all cases, regardless of the operating point or the fusion rule considered. However, the Efficiency of the algorithm varies, and from that variation some criteria for choosing a fusion rule for the multimodal system were inferred.

Even though the results presented here are based on simple fusion rules, the experimental framework can be easily extended to more complex architectures. Future work considering other biometric modalities and fusion schemes will be carried out in order to reach a deeper understanding of the behaviour of multimodal biometric systems under indirect attacks.

Works such as the one presented here emphasize the importance of developing appropriate template protection countermeasures that minimize the effects of

the studied attacks. Some countermeasures have been proposed to counterfeit spoofing attacks, such as [19]. However, the application of those measures against indirect attacks is not straightforward, since they work on raw biometric traits instead of preprocessed templates.

Acknowledgements. This work has been partially supported by projects Contexts (S2009/TIC-1485) from CAM, Bio-Challenge (TEC2009-11186) and Bio-Shield (TEC2012-34881) from Spanish MINECO, TABULA RASA (FP7-ICT-257289) and BEAT (FP7-SEC-284989) from EU, and *Cátedra UAM-Telefónica*. Marta Gomez-Barrero is supported by a FPU Fellowship from Spanish MECED.

References

1. Jain, A.K., et al.: Biometrics: a tool for information security. *IEEE TIFS* 1(2), 125–143 (2006)
2. Schneier, B.: Inside risks: the uses and abuses of biometrics. *Commun. ACM* 42, 136 (1999)
3. Galbally, J., et al.: Evaluation of direct attacks to fingerprint verification systems. *Telecommunication Systems* 47, 243–254 (2011)
4. Galbally, J., et al.: On the vulnerability of face verification systems to hill-climbing attacks. *PR* 43, 1027–1038 (2010)
5. Akhtar, Z., et al.: Robustness analysis of likelihood ratio score fusion rule for multimodal biometric systems under spoof attacks. In: *Proc. ICCST*, pp. 1–8 (2011)
6. TABULA RASA: Trusted biometrics under spoofing attacks (2013)
7. BEAT: Biometrics evaluation and testing (2013)
8. Rodrigues, R., et al.: Evaluation of biometric spoofing in a multimodal system. In: *Proc. IEEE BTAS* (September 2010)
9. Johnson, P.A., et al.: Multimodal fusion vulnerability to non-zero effort (spoof) imposters. In: *Proc. WIFS* (2010)
10. Akhtar, Z., et al.: Spoof attacks in multimodal biometric systems. In: *Proc. IPCSIT*, vol. 4, pp. 46–51. IACSIT Press (2011)
11. Gomez-Barrero, M., et al.: Efficient software attack to multimodal biometric systems and its application to face and iris fusion. *PRL* (2013), doi:10.1016/j.patrec.2013.04.029
12. Masek, L., Kovesi, P.: Matlab source code for a biometric identification system based on iris patterns. Master's thesis, University of Western Australia (2003)
13. Phillips, J., et al.: Overview of the face recognition grand challenge. In: *Proc. IEEE CVPR*, pp. 947–954 (2005)
14. Jain, A.K., et al.: Score normalization in multimodal biometric systems. *PR* 38, 2270–2285 (2005)
15. Kittler, J., et al.: On combining classifiers. *IEEE TPAMI* 20(3), 226–239 (1998)
16. Fierrez, J.: Adapted Fusion Schemes for Multimodal Biometric Authentication. PhD thesis, Universidad Politecnica de Madrid (2006)
17. Ortega-Garcia, J., et al.: The multi-scenario multi-environment BioSecure multimodal database (BMDB). *IEEE TPAMI* 32, 1097–1111 (2010)
18. ANSI-X9.84-2001: Biometric information management and security (2001)
19. Marfella, L., et al.: Liveness-based fusion approaches in multibiometrics. In: *Proc. IEEE BIOMS* (2012)