



Repositorio Institucional de la Universidad Autónoma de Madrid

<https://repositorio.uam.es>

Esta es la **versión de autor** del artículo publicado en:

This is an **author produced version** of a paper published in:

Pattern Recognition Letters Vol. 36 (2014): 243 – 253

DOI: <http://dx.doi.org/10.1016/j.patrec.2013.04.029>

Copyright: © 2014 Elsevier B.V.

El acceso a la versión del editor puede requerir la suscripción del recurso

Access to the published version may require subscription

Efficient Software Attack to Multimodal Biometric Systems and its Application to Face and Iris Fusion

Marta Gomez-Barrero^{b,1,*}, Javier Galbally^{b,1}, Julian Fierrez^{b,1}

*^aBiometric Recognition Group-ATVS, EPS, Universidad Autonoma de Madrid,
C/ Francisco Tomas y Valiente 11, 28049 Madrid, Spain.*

*Corresponding author: Marta Gomez-Barrero, e-mail: marta.barrero@uam.es, tel.: +34 91 497 33 63

Email addresses: marta.barrero@uam.es (Marta Gomez-Barrero),
javier.galbally@uam.es (Javier Galbally), julian.fierrez@uam.es
(Julian Fierrez)

¹This 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*.

Efficient Software Attack to Multimodal Biometric Systems and its Application to Face and Iris Fusion

Marta Gomez-Barrero^{b,1,*}, Javier Galbally^{b,1}, Julian Fierrez^{b,1}

^b*Biometric Recognition Group-ATVS, EPS, Universidad Autonoma de Madrid,
C/ Francisco Tomas y Valiente 11, 28049 Madrid, Spain.*

Abstract

In certain applications based on multimodal interaction it may be crucial to determine not only *what* the user is doing (commands), but *who* is doing it, in order to prevent fraudulent use of the system. The biometric technology, and particularly the multimodal biometric systems, represent a highly efficient automatic recognition solution for this type of applications.

Although multimodal biometric systems have been traditionally regarded as more secure than unimodal systems, their vulnerabilities to spoofing attacks have been recently shown. New fusion techniques have been proposed and their performance thoroughly analysed in an attempt to increase the robustness of multimodal systems to these spoofing attacks. However, the vulnerabilities of multimodal

*Corresponding author: Marta Gomez-Barrero, e-mail: marta.barrero@uam.es, tel.: +34 91 497 33 63

Email addresses: marta.barrero@uam.es (Marta Gomez-Barrero),
javier.galbally@uam.es (Javier Galbally), julian.fierrez@uam.es
(Julian Fierrez)

¹This 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*.

approaches to software-based attacks still remain unexplored. In this work we present the first software attack against multimodal biometric systems. Its performance is tested against a multimodal system based on face and iris, showing the vulnerabilities of the system to this new type of threat. Score quantization is afterwards studied as a possible countermeasure, managing to cancel the effects of the proposed attacking methodology under certain scenarios.

Keywords:

Multimodal system, security, vulnerabilities, hill-climbing, countermeasures.

1. Introduction

Multimodal systems represent a new direction for computing that embraces users' natural behaviour as the center of human-computer interaction [1]. As with any other novel discipline, the research community is just beginning to understand how to design robust and well integrated multimodal systems. But only through multidisciplinary cooperation among those with expertise in individual component technologies can multimodal systems reach its final aim: building more general and robust systems that will reshape daily computing tasks and have significant commercial impact [2].

One of the main areas of research in multimodal interaction, where specific expertise is needed, is *recognition*, generally regarded as a form of processing users' commands. However, for certain applications based on multimodal interaction, a second form of recognition is crucial: it is not only necessary to distinguish *what* the user is doing, but *who* is doing it, so that non-authorized individuals cannot use the system. For these cases, a robust personal automatic recognition solution such as the one provided by *biometrics* is required. Although being relatively young compared to other mature and long-used security technologies, biometrics have

18 emerged in the last decade as a pushing alternative for applications where auto-
19 matic recognition of people is needed. Certainly, biometrics are very attractive
20 and useful for the final user: forget about PINs and passwords, you are your own
21 key [3, 4]. However, we cannot forget that as any technology aimed to provide
22 a security service, biometric systems are exposed to external attacks which could
23 compromise their integrity [5]. Thus, it is of special relevance to understand the
24 threats to which they are subjected and to analyse their vulnerabilities in order to
25 prevent possible attacks and increase their benefits for the users.

26 External attacks to biometric systems are commonly divided into: *direct at-*
27 *tacks* (also known as *spoofing attacks*), carried out against the sensor, and *indirect*
28 *attacks*, directed to some of the inner modules of the system. In the last recent
29 years important research efforts have been conducted to study the vulnerabilities
30 of biometric systems to both direct and indirect attacks [6, 7, 8, 9].

31 This new concern which has arisen in the biometric community regarding the
32 security of biometric systems has led to the appearance of several international
33 projects, like the European Tabula Rasa [10], which base their research on the
34 security through transparency principle [11, 12]: in order to make biometric sys-
35 tems more secure and reliable, their vulnerabilities need to be analysed and useful
36 countermeasures need to be developed.

37 In this scenario, biometric multimodality has been regarded as an effective
38 way of increasing the robustness of biometric-based security systems to external
39 attacks. Combining the information offered by several traits would force an even-
40 tual intruder to successfully break several unimodal modules instead of just one.
41 However, it has already been proven that this is not necessary in spoofing attacks:
42 breaking into the module based on the most accurate biometric trait grants access

43 to the multimodal system in many occasions [13, 14, 15].

44 In addition to research works which address the vulnerabilities of multimodal
45 systems to spoofing attacks [13, 14, 15, 16, 17, 18, 19, 20], different studies
46 may be found in the literature regarding the analysis of indirect attacks against
47 unimodal systems [8, 9, 21]. However, the problem of whether multimodal ap-
48 proaches are vulnerable or not to software-based attacking methodologies still
49 remains unexplored.

50 In the present work we propose and analyse a general multimodal indirect at-
51 tack, which can be used to study the vulnerabilities of biometric systems based on
52 different number of traits, different fusion strategies and different types of tem-
53 plates (e.g., real valued, binary). Without loss of generality, the attack is applied
54 to the particular case of a face- and iris-based recognition system. This trait com-
55 bination is regarded as one of the most popular and user-friendly, since the acqui-
56 sition of both traits can be transparent to the user [22, 23, 24, 25]. This provides
57 a straight-forward integration of both modalities, a complex topic on multimodal
58 computation [26]. Furthermore, the experimental protocol used is fully replicable,
59 so that the results obtained can be fairly compared.

60 Score quantization is studied afterwards as a possible countermeasure against
61 the proposed attack. Two different approaches are analysed: quantizing the score
62 before and after the fusion of the partial face and iris scores. While the second
63 scheme barely reduces the success rate and efficiency of the attack, the first one
64 succeeds in preventing an intruder from breaking into the system.

65 Thus, following the same transparency principle which is starting to prevail in
66 the biometric community through European Projects such as Tabula Rasa [11, 12],
67 the main objectives and contributions of the present work are: *i*) proposal of a

68 fully novel software-based attacking methodology against multimodal systems,
69 *ii*) study of the vulnerabilities of a realistic multimodal system to the previous at-
70 tack under a replicable scenario, *iii*) comparison of the performance of the attack
71 to that obtained against the unimodal modules in order to determine if the mul-
72 timodal approach increases the security of the system against this type of threat,
73 and *iv*) study of some biometric-based countermeasures which may prevent such
74 an attack.

75 The paper is structured as follows. Related works are summarised in Sect. 2.
76 The novel multimodal attacking algorithm used to evaluate the system is presented
77 in Sect. 3. Then the multimodal verification system evaluated is described in
78 Sect. 4. The database and experimental protocol followed are presented in Sect. 5.
79 In Sect. 6 we describe and analyse the results obtained. Score quantization is
80 studied as a possible countermeasure in Sect. 7. Conclusions are finally drawn in
81 Sect. 8.

82 2. Related Works

83 In 2001, Ratha *et al.* identified and classified in a biometric recognition sys-
84 tem eight possible points of attack [27]. These vulnerable points can be broadly
85 divided into direct and indirect attacks.

86 **Direct attacks.** Also known as spoofing-attacks, these are attacks at the sensor
87 level, carried out with synthetic biometric traits, such as gummy fingers or high
88 quality printed iris images, and thus requiring no knowledge for the attacker of
89 the inner parts of the system (matching algorithm used, feature extraction method,
90 template format, etc.) Some research regarding the vulnerabilities of multimodal
91 systems to these attacks has been carried out over the last recent years: in 2005,

Chetty and Wagner [14] tested the performance of spoofing attacks against a novel multimodal system based on face and voice; in 2009, Tan [28] investigated methods for increasing the security of multimodal systems based on face and voice against spoofing attacks; in 2010 [16] and 2011 [15], Rodrigues *et al.* evaluated the vulnerabilities of a multimodal system based on face and fingerprint, using different fusion techniques and proposing new ones; in 2010, Johnson *et al.* [19] analysed the effect of spoofing attacks against a multimodal system based on face and iris, proposing a method for the vulnerabilities assessment of these systems; later in 2010, Marasco [20] analysed the security risks in multimodal biometric systems based on face and fingerprint coming from spoofing attacks; in 2011, Akthar *et al.* [13, 17] used real rather than simulated spoof samples for the evaluation of the vulnerabilities of a multimodal system based on fingerprint, face and iris, proposing a new learning algorithm able to improve the security offered by the system against spoofing attacks. All these works have proven that combining several traits in one system for person authentication does not necessarily increment the security offered against spoofing attacks, since the system can be bypassed by breaking only one of the unimodal traits.

Indirect attacks. These attacks are directed to the inner modules of the system and can be further divided into three groups, namely: *i*) attacks to the communication channels between modules of the system, extracting, adding or changing information; *ii*) attacks to the feature extractor and the matcher may be carried out using a Trojan Horse that bypasses the corresponding module; and *iii*) attacks to the system database which manipulate it in order to gain access to the application, by changing, adding or deleting a template. While for direct attacks the intruder needed no knowledge about the inner modules of the system, this

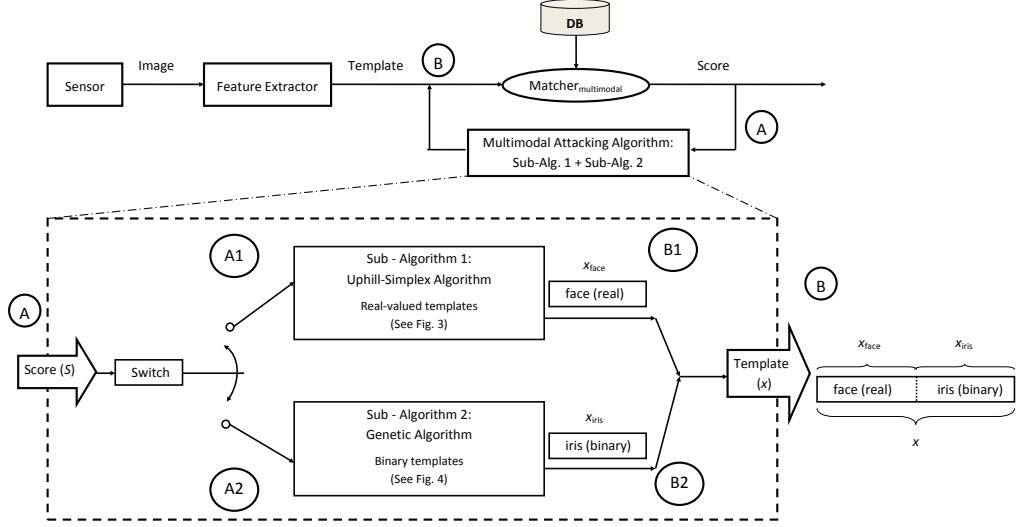


Figure 1: Diagram of a general hill-climbing attack (top), with the specific modification scheme for the combined algorithm (bottom).

117 knowledge is a main requisite here, together with access to some of the system
 118 components (database, feature extractor, matcher, etc.). Most of these indirect
 119 attacks are based on some variation of a hill-climbing algorithm, consisting on
 120 iteratively changing some synthetically generated templates until access to the
 121 system is granted. Even though some research has been done in this area using
 122 unimodal systems [8, 9, 21, 29], to the best of our knowledge there is no previ-
 123 ous analysis of the vulnerabilities of multimodal biometric systems to this kind of
 124 attacks.

125 3. Proposed Attack

126 Until now, only the vulnerabilities of unimodal systems to indirect attacks have
 127 been analysed. In this section we present the first algorithm for the evaluation of

the vulnerabilities of multimodal systems to this type of threat. As can be observed in Fig. 1 (top), the input to the algorithm are the scores given by the matcher, and the output the templates to be compared to the client account.

For simplicity, the attacking methodology is described here for the particular case of a multimodal system based on the score fusion of a real valued (e.g. face) and a binary (e.g. iris) matcher. However, the proposed approach is general and may be applied with very small modifications to attack multimodal systems working on: *i*) more than two traits represented with real-valued or binary templates (by adding new blocks after the switch in Fig. 1), or *ii*) feature-based fusion strategies (by rearranging the template disposition).

In order to attack a multimodal biometric system where one of the biometric traits is represented with real values and the other is binary (most iris recognition systems work on binary templates), the algorithm here presented combines two sub-algorithms. Each of them attacks one segment of the template: the real-valued or the binary segment. In the following subsections, each of the individual sub-algorithms is described. Finally, the multimodal attacking algorithm based on the previous two models is presented.

3.1. Sub-Algorithm 1: Hill-Climbing based on the Uphill Simplex Algorithm

Problem statement. Consider the problem of finding a K -dimensional vector of real values x_{face} which, compared to an unknown template $\mathcal{C}_{\text{face}}$ (in our case related to a specific client), produces a similarity score bigger than a certain threshold δ_{face} , according to some matching function J_{face} , i.e., $J_{\text{face}}(\mathcal{C}_{\text{face}}, x_{\text{face}}) > \delta_{\text{face}}$. The template can be another K -dimensional vector or a generative model of K -dimensional vectors.

Assumptions. Let us assume:

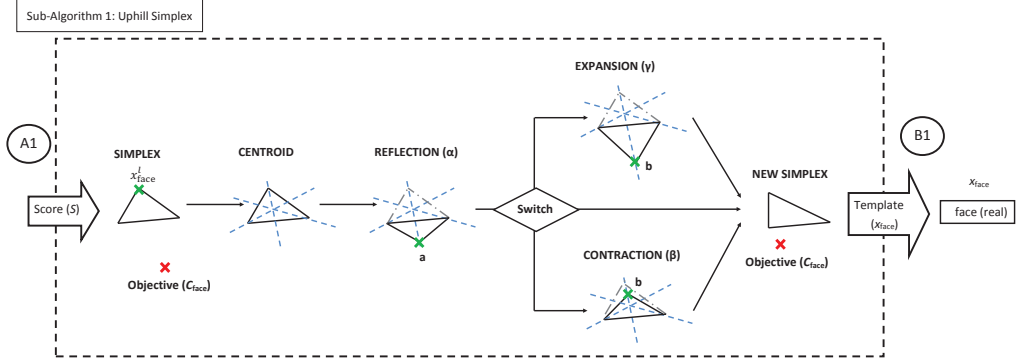


Figure 2: Diagram of the modification scheme for the Sub-Algorithm 1, based on the Uphill-Simplex.

- That there exists a statistical model G (K -variate Gaussian with mean μ_G and a diagonal covariance matrix Σ_G , with $\sigma_G^2 = \text{diag}(\Sigma_G)$), in our case related to a background set of users, overlapping to some extent with $\mathcal{C}_{\text{face}}$.
- That we have access to the evaluation of the matching function $J_{\text{face}}(\mathcal{C}_{\text{face}}, x_{\text{face}})$ for several trials of x_{face} .

Algorithm. The problem stated above can be solved by adapting the Downhill Simplex algorithm first presented in [30] to maximize instead of minimize the function J_{face} . We iteratively form new simplices by reflecting one point, x_{face}^l , in the hyperplane of the remaining points, until we are close enough to the maximum of the function. The point to be reflected will always be the one with the lowest value given by the matching function, since it is in principle the one furthest from our objective. Thus, as can be observed in Fig. 2, the different steps followed by the sub-algorithm 1 are:

1. Compute the statistical model $G(\mu_G, \sigma_G)$ from a development pool of users.

- 167 2. Take $K + 1$ samples (x_{face}^i) defining the initial simplex from the statistical
 168 model G and compute the similarity scores $J_{\text{face}}(\mathcal{C}_{\text{face}}, x_{\text{face}}^i) = s_{\text{face}}^i$, with
 169 $i = 1, \dots, K + 1$.
- 170 3. Compute the centroid \bar{x}_{face} of the simplex as the average of x_{face}^i : $\bar{x}_{\text{face}} =$
 171 $\frac{1}{K+1} \sum_i x_{\text{face}}^i$.
- 172 4. Reflect the point x_{face}^l according to the next steps, adapted from the Down-
 173 hill Simplex algorithm [30]. In the following, the indices l and h are defined
 174 as $h = \arg \max_i (s_{\text{face}}^i)$, $l = \arg \min_i (s_{\text{face}}^i)$.

4. a) **Reflection:** Given a constant $\alpha > 0$, the *reflection coefficient*, we compute:

$$a = (1 + \alpha)\bar{x}_{\text{face}} - \alpha x_{\text{face}}^l.$$

175 Thus, a is on the line between x_{face}^l and \bar{x}_{face} being α the ratio between
 176 the distances $[a\bar{x}_{\text{face}}]$ and $[x_{\text{face}}^l\bar{x}_{\text{face}}]$. If $s_{\text{face}}^l < s_{\text{face}}^a < s_{\text{face}}^h$ we replace
 177 x_{face}^l by a . Otherwise, we go on to step 4b.

- 178 4. b) **Expansion or contraction.**

- i. **Expansion:** If $s_{\text{face}}^a > s_{\text{face}}^h$ (i.e., we have a new maximum) we expand a to b as follows:

$$b = \gamma a + (1 - \gamma)\bar{x}_{\text{face}},$$

179 where $\gamma > 1$ is another constant called *expansion coefficient*,
 180 which represents the ratio between the distances $[b\bar{x}_{\text{face}}]$ and $[s\bar{x}_{\text{face}}]$.
 181 If $s_{\text{face}}^b > s_{\text{face}}^h$, we replace x_{face}^l by b . Otherwise, we have a failed
 182 expansion and replace x_{face}^l by a .

- ii. **Contraction:** If we have reached this step, then $s_{\text{face}}^a \leq s_{\text{face}}^l$ (i.e. replacing x_{face}^l by a would leave s_{face}^a as the new minimum). We

compute

$$b = \beta x_{\text{face}}^l + (1 - \beta) \bar{x}_{\text{face}},$$

183 where $0 < \beta < 1$ is the *contraction coefficient*, defined as the
 184 ratio between the distances $[b\bar{x}_{\text{face}}]$ and $[x_{\text{face}}^l\bar{x}_{\text{face}}]$. If $s_{\text{face}}^b >$
 185 $\max(s_{\text{face}}^l, s_{\text{face}}^a)$, then we replace x_{face}^l by b ; otherwise, the con-
 186 tracted point is worse than x_{face}^l , and for such a failed contraction
 187 we replace all the x_{face}^i 's by $(x_{\text{face}}^i + x_{\text{face}}^h)/2$.

188 5. With the new x_{face}^l value, update the simplex and return to step 3.

189 **Stopping criteria.** The algorithm stops when: *i*) the maximum similarity
 190 score of the simplex vertices is higher than the threshold δ_{face} (i.e., the account is
 191 broken), *ii*) the variation of the similarity scores obtained in a number of itera-
 192 tions is lower than a certain threshold or *iii*) a maximum number of iterations is
 193 reached.

194 **Additional note.** It is important to notice for the computation of the Efficiency
 195 (defined in Sect. 5.3) of this sub-algorithm that at each iteration (except for the
 196 initial one) a maximum of 2 matchings will be performed (i.e., $s_{\text{face}}^a + s_{\text{face}}^b$). On
 197 average, the number of matchings computed per iteration will be lower than 2 and
 198 greater than 1.

199 The hill-climbing based on the Uphill Simplex algorithm was first presented in
 200 [31], where it was used to successfully attack a signature verification system. The
 201 performance of the proposed algorithm showed a clear improvement in the attack-
 202 ing capabilities with respect to previously proposed state-of-the-art approaches,
 203 which motivated its choice for the present multimodal vulnerability study.

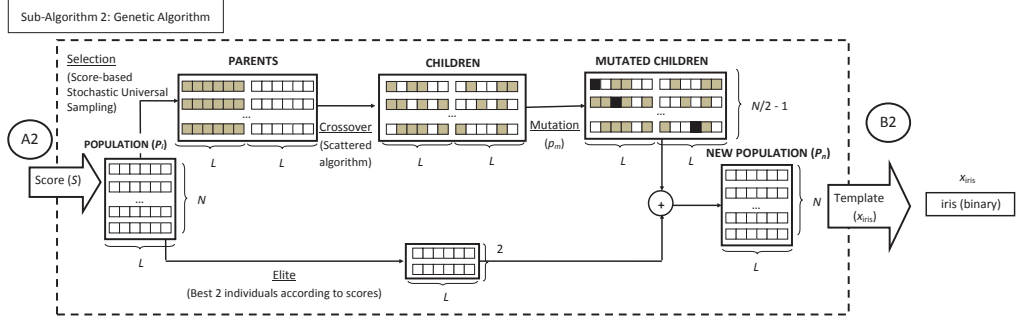


Figure 3: Diagram of the modification scheme for the Sub-Algorithm 2, based on a genetic algorithm.

3.2. Sub-Algorithm 2: Indirect Attack based on a Genetic Algorithm

Problem statement. Consider the problem of finding an L -dimensional binary vector x_{iris} which, compared to an unknown template $\mathcal{C}_{\text{iris}}$ (in our case related to a specific client), produces a similarity score bigger than a certain threshold δ_{iris} , according to some matching function J_{iris} , i.e., $J_{\text{iris}}(\mathcal{C}_{\text{iris}}, x_{\text{iris}}) > \delta_{\text{iris}}$. The template can be another L -dimensional vector or a generative model of L -dimensional vectors.

Assumptions. Let us assume:

- That we have access to the evaluation of the matching function $J_{\text{iris}}(\mathcal{C}_{\text{iris}}, x_{\text{iris}})$ for several trials of x_{iris} .

Algorithm. The problem stated above may be solved by using a genetic algorithm, which has shown a remarkable performance in binary optimization problems [32], to optimize the similarity score given by the matcher, that is, the fitness value for an individual is $s_{\text{iris}} = \mathcal{J}_{\text{iris}}(x_{\text{iris}}, \mathcal{C}_{\text{iris}})$. As can be seen in Fig. 3 the steps followed by the sub-algorithm 2 are:

- 219 1. Generate an initial population P_i with N individuals of length L , L being
220 the length of the iris code.
- 221 2. Compute the similarity scores s^i of the individuals (x_{iris}^i) of the population
222 P_i , $s_i = J(x_{\text{iris}}^i, \mathcal{C}_{\text{iris}})$ with $i = 1, \dots, N$.
- 223 3. Four rules are used at each iteration to create the next generation P_n of
224 individuals from the current population:
 - 225 3. a) **Elite**: the two individuals with the maximum similarity scores are kept
226 unaltered for the next generation.
 - 227 3. b) **Selection**: certain individuals, the *parents*, are chosen by stochastic
228 universal sampling [33]. This way, the individuals with the highest fit-
229 ness values (similarity scores) are more likely to be chosen as parents
230 for the next generation: one subject can be selected from 0 to many
231 times. From the original N individuals, $N/2 - 1$ *fathers* and $N/2 - 1$
232 *mothers* are chosen.
 - 233 3. c) **Crossover**: parents are combined to form the $N - 2$ *children* of the
234 next generation, following a scattered crossover method. A random
235 binary vector is created and the genes (bits) of the child are selected
236 from the first parent where the value of the random vector is 1, and
237 from the second when it is 0 (vice versa for the second child).
 - 238 3. d) **Mutation**: random changes are applied to the bit values of the new
239 children with a mutation probability p_m .
- 240 4. Redefine $P_i = P_n$ and return to step 2.

241 **Stopping criteria.** The algorithm stops when: *i*) the best fitness score is
242 higher than the threshold δ_{iris} (i.e., the account is broken), *ii*) the variation of the

243 similarity scores obtained in a number of generations is lower than a previously
244 fixed value, or *iii*) when the maximum number of generations is reached.

245 **Additional note.** It is important to notice for the computation of the Efficiency
246 (defined in Sect. 5.3) of this sub-algorithm that at each iteration (i.e., generation)
247 N matchings are performed (one for each of the members of the population).

248 This particular implementation of a genetic algorithm was first presented in
249 [34], where it was used to analyse the vulnerabilities of the same iris recogni-
250 tion system considered in this work. The performance of the proposed algorithm
251 showed a very high attacking potential with very encouraging results and was the
252 first one, to our knowledge, working on a binary input (such as the iriscodes).
253 Therefore, its use as part of the global multimodal attack presented here seemed
254 like a promising choice.

255 3.3. Multimodal Attack: Combination of Sub-Algorithms 1 (Uphill-Simplex) and 256 2 (Genetic-Algorithm)

257 **Problem statement.** Consider the problem of finding a $(K + L)$ -dimensional
258 vector x of real and binary values which, compared to an unknown template \mathcal{C}
259 (in our case related to a specific client), produces a similarity score bigger than
260 a certain threshold δ , according to some matching function J , i.e., $J(\mathcal{C}, x) > \delta$.
261 The template can be another $(K + L)$ -dimensional vector or a generative model
262 of $(K + L)$ -dimensional vectors.

263 **Assumptions.** Let us assume:

- 264 • That we know the distribution of the two subtemplates (real-valued x_{face}
265 and binary x_{iris}) within the multimodal template x .

- That we have access to the evaluation of the matching function $J(\mathcal{C}, x)$ for several trials of x .

Algorithm. The problem stated above may be solved by dividing the template x into its real-valued (x_{face}) and binary parts (x_{iris}) and alternately optimize each of them as can be seen in Fig. 1. In order to optimize each of the parts, the algorithms described in the previous subsections are used: the Sub-Algorithm 1 for the real-valued segment (face) and the Sub-Algorithm 2 for the binary segment (iris). Thus, the steps followed are:

1. Generate a synthetic template (x) randomly initializing the real-valued (x_{face}) and binary (x_{iris}) segments, and compute the similarity score $S = J(\mathcal{C}, x)$, which will be used as optimization criterion.
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.

Stopping criteria. 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 (i.e., changes between the optimized segments) exceeds a previously fixed value (i.e., the attack has failed).

Additional note. As will be analysed in the experimental section this algorithm may present different results depending on whether it starts attacking the real-valued or binary part of the template.

289 It is very important to notice that the multimodal attacking algorithm does not
290 have access at any point to the partial scores of the unimodal modules (s_{face} and
291 s_{iris}) but only uses the final fused score given by the system (S). This way, in the
292 description of the previous two sub-algorithms, s_{face} and s_{iris} should be changed by
293 S when they are used as part of the multimodal attack and not individually.

294 Both attacking sub-algorithms stop when the improvement of the final multi-
295 modal score saturates (i.e., the variation of the multimodal similarity scores ob-
296 tained in a number of iterations or generations is lower than a certain threshold).
297 This “switching” methodology is preferred over a “sequential” approach based on
298 the assumption that once the algorithm has saturated attacking one of the unimodal
299 subsystems, further changes in the other modality will lead to new improvements
300 in the final multimodal score.

301 **4. Multimodal Verification System Attacked**

302 The multimodal verification system evaluated in this work is the fusion of two
303 unimodal systems, namely: *i*) a modified version of the iris recognition system
304 developed by L. Masek² [35], which is widely used in many iris related publica-
305 tions; and *ii*) an Eigenface-based face verification system [36], used to present
306 initial face verification results for the recent Face Recognition Grand Challenge
307 [37].

308 *4.1. Face Verification System*

309 The system evaluated uses Multi-Layer Perceptron (MLP) and a cascade of
310 Haar-like classifiers in order to segment the faces in the images, together with

²The source can be freely downloaded from www.csse.uwa.edu.au/pk/studentprojects/libor/sourcecode.html

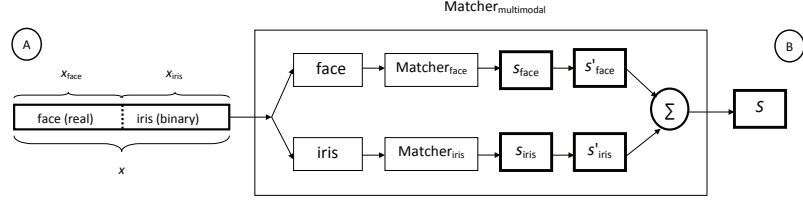


Figure 4: Similarity score obtained from one multimodal template (x) consisting of two different segments, containing: face features (x_{face} , real values) and the iris code (x_{iris} , binary). The unimodal verification subsystems give the corresponding scores (s_{face} , s_{iris}), which are then normalised (s'_{face} , s'_{iris}) and fused to obtain the final output of the global system: S .

the position of the eyes on them. Principal Component Analysis (PCA) is used afterwards so that face images can be represented in a lower dimensional space [8]. 80% of the variance is retained when training the PCA vector space with cropped face images of size 64×80 , reducing the original 5120-dimensional space to only 100 dimensions or eigenvectors.

Finally, the similarity scores are computed in this PCA vector space using the Euclidean distance.

4.2. Iris Verification System

The system comprises four different steps, namely: *i) segmentation*, where the method proposed in [38] is followed, modelling the iris and pupil boundaries as circles; *ii) normalization*, using a technique based on Daugman's ruber sheet model that maps the segmented iris region into a 2D array [39]; *iii) feature encoding*, which produces a binary template of $20 \times 480 = 9,600$ bits and the corresponding noise mark (representing the eyelids areas) by convolving the normalized iris pattern with 1D Log-Gabor wavelets; and *iv) matching*, where the inverse of a modified Hamming distance is used, which takes into account the

327 noise mask bits.

This way, the similarity score between two templates is computed as $1/HD$ (so that a higher score implies a higher degree of similarity):

$$HD = \frac{\sum_{j=1}^L X_j(XOR)Y_j(AND)\bar{X}n_j(AND)\bar{Y}n_j}{L - \sum_{k=1}^L Xn_k(OR)Yn_k}$$

328 where X_j and Y_j are the two bitwise templates to compare, Xn_j and Yn_j are the
329 corresponding noise masks for X_j and Y_j , and L is the number of bits comprised
330 in each template. $\bar{X}n_j$ denotes the logical not operation applied to Xn_j .

331 4.3. Multimodal Verification System

332 Given an input vector x , the system performs the following tasks in order to
333 obtain the final score, S , as can be seen in Fig. 4:

- 334 1. Compute the similarity scores obtained by the face (s_{face}) and iris (s_{iris})
335 traits, as given by the matchers described in Sect. 4.1 and Sect. 4.2.
2. Normalize the scores s_k , with $k = \{\text{face}, \text{iris}\}$, using hyperbolic tangent
estimators (its robustness and high efficiency are proven in [40]):

$$s'_k = \frac{1}{2} \left\{ \tanh \left(0.01 \frac{s_k - \mu}{\sigma} \right) + 1 \right\}$$

336 where s_k is the original similarity score obtained by the iris (respectively
337 face) section of the template, μ and σ the mean and standard deviation of
338 the scores distribution of the iris (respectively face), and s'_k the normalised
339 score. This way, both partial scores (face and iris) lie in the interval $[0, 1]$.

3. Finally, both normalised scores are fused with a sum, given the very good
results that this fusion rule has presented even when compared with more
sophisticated methods like decision trees [41] or neural networks [22]:

$$S = s'_{\text{iris}} + s'_{\text{face}}$$

340 There may be other fusion strategies that can improve the performance of
341 the multimodal system. However, simple summation gives very good re-
342 sults, and it is not the scope of the paper to find the optimal fusion strategy.

343 **5. Database and Experimental Protocol**

344 Prior to any vulnerability assessment study a performance evaluation of the
345 systems being attacked should be carried out. The performance evaluation will
346 permit to determine how good is the system and, more important, the operating
347 points where it will be attacked as the success chances of this kind of attacking
348 algorithms are, in principle, highly dependent on the False Acceptance and False
349 Rejection rates of the system. While the FRR measures the probability of rejecting
350 a genuine user, the FAR gives a measure of the probability of an impostor being
351 taken as a genuine user. Therefore, in general, the higher the FAR, the easier for an
352 eventual attacker to break into the system. Moreover, for the particular case of the
353 proposed method, attacking the system at a lower FAR implies reaching a higher
354 threshold, which leads to a decrease on the success chances of the algorithm.

355 Furthermore, defining the operating points will enable us to compare, in a
356 more fair fashion, the vulnerabilities of the different systems to the same attack
357 (i.e., we can determine for a given FAR or FRR which of them is less/more robust
358 to the attacking approach).

359 Both the database and the protocol used for the performance and security eval-
360 uations of the multimodal system are the same ones used for the evaluation of
361 the unimodal subsystems, so that the results are fully comparable. This way, we
362 will be able to determine whether the multimodality enhances the system security
363 against the proposed attacking approaches with respect to the unimodal traits.

364 5.1. Database

365 The experiments are carried out on the face and iris subcorpora included in
366 the Desktop Dataset of the multimodal BioSecure database [42], which comprises
367 voice, fingerprints, face, iris, signature and hand of 210 users, captured in two
368 time-spaced acquisition sessions. This database was acquired thanks to the joint
369 effort of 11 European institutions and has become one of the standard benchmarks
370 for biometric performance and security evaluations [43]. It is publicly available
371 through the BioSecure Foundation³.

372 The database comprises three datasets captured under different acquisition
373 scenarios, namely: *i*) Internet Dataset (DS1, captured through the Internet in an
374 unsupervised setup), *ii*) Desktop Dataset (DS2, captured in an office-like envi-
375 ronment with human supervision), and *iii*) the Mobile Dataset (DS3, acquired on
376 mobile devices with uncontrolled conditions). The face subset used in this work
377 includes four frontal images (two per session) with an homogeneous grey back-
378 ground, and captured with a reflex digital camera without flash ($210 \times 4 = 840$
379 face samples), while the iris subset includes four grey-scale images (two per ses-
380 sion as well) per eye, all captured with the Iris Access EOU3000 sensor from LG.
381 In the experiments only the right eye of each user has been considered, leading
382 this way as in the face case to $210 \times 4 = 840$ iris samples.

383 5.2. Performance evaluation

384 As the iris and face subcorpus present identical sample distributions, the pro-
385 tocol followed for the performance evaluation of the unimodal modules and the
386 multimodal system is the same. As can be seen in Fig. 5, each subcorpus of the

³<http://biosecure.it-sudparis.eu/AB>

		BioSecure DS2 DB (210 Users)	
Session	Sample	170 Users	40 Users
1	1	Training	Test (Impostors)
	2		
2	1	Test (Clients)	
	2		

Figure 5: Partition of the BioSecure DS2 DB according to the performance evaluation protocol defined.

387 database is divided in two sets, namely: *i*) a training set comprising the first three
388 samples of 170 clients, used as the enrolment templates; *ii*) a test set formed by
389 the fourth image of the 170 clients above (used to compute the genuine scores)
390 and the 4 images of the remaining 40 users (used to compute the impostor scores).

391 As a result of: *i*) using the same subjects for PCA training and client enrol-
392 ment for the face verification subsystem, and *ii*) manually segmenting those eyes
393 that were not successfully segmented automatically (3.04%), the system perfor-
394 mance is optimistically biased, and therefore harder to attack than in a practical
395 situation (in which the enrolled clients may not have been used for PCA training
396 and the image segmentation would be fully automatic). This means that the results
397 presented in this paper are a conservative estimate of the attack’s performance.

398 The final score given by the system is the average of the scores obtained after
399 matching the input template to the three face and iris templates of the client model
400 \mathcal{C} . Table 1 shows that the ERR of the unimodal face and iris modules and of the
401 whole multimodal system computed according to the protocol described above.
402 In this chart we can observe that: *i*) the performance of the unimodal modules is
403 not noticeably affected by score normalization (i.e., the EER barely changes after
404 normalising the scores), and *ii*) the performance of the multimodal system (0.83%

Table 1: EER of the unimodal and multimodal systems, based on face and iris, before and after the normalization of the scores.

	EER (%)		
	Face	Iris	Multimodal
Before Norm.	6.55	4.11	-
After Norm.	6.61	4.04	0.83

405 EER) clearly improves that of the unimodal systems (4% and 6% respectively). In
 406 Fig. 6 the Detection Error Tradeoff (DET) curves of the unimodal and multimodal
 407 systems obtained using the described protocol are shown. As can be seen, the
 408 multimodal system clearly outperforms both unimodal systems at all points.

409 5.3. Experimental Protocol for the Attacks

410 The user accounts to be attacked by the algorithm are generated with the train-
 411 ing set defined in the performance evaluation protocol (i.e., the first three sam-
 412 ples of the 170 users in Fig. 5). The performance of the attack is evaluated in
 413 terms of: *i*) its Success Rate (SR) or expected probability of bypassing the sys-
 414 tem, computed as the ratio $SR = A_B/A_T$, where A_B is the number of broken
 415 accounts and A_T is the total number of attacked accounts; and *ii*) its Efficiency
 416 (Eff), or inverse of the average number of comparisons needed to break an ac-
 417 count, $Eff = 1 / \left(\sum_{i=1}^{A_B} n_i / A_B \right)$, where n_i is the number of comparisons made to
 418 bypass the i th account, with $i = 1, \dots, A_B$.

419 It has to be emphasized that the Eff is computed in terms of the number of
 420 *matchings* performed by the attacking algorithm and not according to the number
 421 of *iterations* needed (i.e., two algorithms performing the same number of itera-

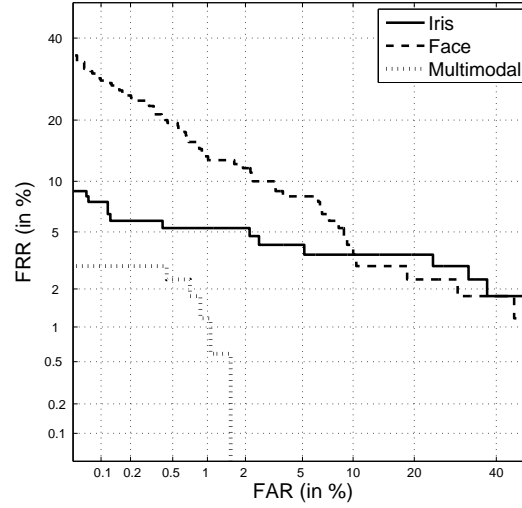


Figure 6: DET curves of the unimodal and multimodal systems.

tions to break an account do not necessarily have the same Eff).

The SR gives an estimation of how dangerous the attack is: the higher the SR, the bigger the threat. On the other hand, the Eff tells us how easy it is for the attack to bypass the system in terms of speed: the higher the Eff, the faster the attack.

The different attacks have been evaluated at three operating points which correspond to $\text{FAR} = 0.1\%$, $\text{FAR} = 0.05\%$ and $\text{FAR} = 0.01\%$, which, according to [44], offer a low, medium and high security level.

6. Results: Attack Performance

The objectives of this first study of the vulnerabilities of a multimodal system to an indirect attack are: *i*) to evaluate the performance of the proposed attacking methodology, and *ii*) to test whether the use of two different biometric traits

Table 2: Eff and SR for the Sub-Algorithm 1 (Uphill-Simplex) and Sub-Algorithm 2 (Genetic Algorithm) attacks carried out against the corresponding unimodal systems, and for the Multimodal Attack against the multimodal system.

FAR	Unimodal Attacks				Multimodal Attack			
	Sub-Alg. 1 vs Face		Sub-Alg. 2 vs Iris		Starts Face		Starts Iris	
	SR	Eff ($\times 10^{-4}$)	SR	Eff ($\times 10^{-4}$)	SR	Eff ($\times 10^{-4}$)	SR	Eff ($\times 10^{-4}$)
0.10%	100%	22.472	91.18%	1.400	100%	1.9372	100%	1.4180
0.05%	100%	22.124	80.89%	1.255	100%	1.8218	100%	1.3585
0.01%	100%	21.930	62.36%	1.102	100%	1.3702	100%	1.1112

434 increments the security level and robustness of the system to this kind of attacks.

435 In the first set of experiments, the performance of the two attacking sub-
 436 algorithms against the unimodal systems is studied, so that later a comparison
 437 between the unimodal and the multimodal systems can be established. In the sec-
 438 ond set, the performance of the attack against the multimodal system is tested.
 439 Score quantization is afterwards analysed as a possible countermeasure, studying
 440 its impact in the SR and the Eff of the multimodal attacking scheme.

441 6.1. Sub-Algorithm 1 vs Face Verification System

442 The performance of the Sub-Algorithm 1 against the unimodal system based
 443 on eigenfaces is tested at the three operating points mentioned before, namely: *i)*
 444 FAR = 0.10%, *ii)* FAR = 0.05%, *iii)* FAR = 0.01%. The results of the experi-
 445 ments are detailed in Table 2, where we can observe that the algorithm success-
 446 fully breaks all the attacked accounts. Also worth noting that for this attack the
 447 efficiency remains almost invariant, regardless of the operating point considered.

448 It should also be emphasized that in the present work the hill-climbing attack
449 is initialized from a normal distribution of zero mean and unit variance, that is,
450 the first simplex is generated without needing any training faces, contrary to what
451 happened in other state of the art attacking methods [8]. Furthermore, the param-
452 eters α , β and γ used here are the same that were optimized in [31] to break a
453 signature verification system, which proves the robustness of the algorithm: it is
454 able to break totally heterogeneous systems working on different biometric traits
455 without adjusting its parameters.

456 6.2. Sub-Algorithm 2 vs Iris Verification System

457 As before, the performance of the Sub-Algorithm 2 against the unimodal sys-
458 tem based on iris is tested at the three operating points mentioned before, namely:
459 *i)* FAR = 0.10%, *ii)* FAR = 0.05%, *iii)* FAR = 0.01%. The results of the experi-
460 ments are also shown in Table 2, where we can observe that the algorithm is able
461 to successfully break more than 90% of the accounts for the point of operation
462 corresponding to a low security level, and more than 60% for the point corre-
463 sponding to a high security level. As in the previous case the efficiency of the
464 attack remains almost invariant, slightly decreasing, as would be expected, for
465 higher security points where the attack needs more iterations to break the system
466 (i.e., it becomes slower).

467 6.3. Combined Attack vs Multimodal System

468 We run two sets of experiments, namely: *i)* the algorithm starts attacking
469 the face section of the template (Sub-Algorithm 1), and *ii)* the algorithm starts
470 attacking the iris section (Sub-Algorithm 2). Between 40% and 60% of the times
471 that the algorithm starts attacking the iris section of the template it is able to break

472 the account without changing to the face segment. This does not happen when the
473 algorithm starts attacking the face segment. This way, as it was already proven for
474 spoofing attacks in [13, 15, 19], attacking only the best individual matcher (i.e.,
475 the unimodal system with the lowest EER, the iris one in our case) grants in many
476 cases access to the system under attack.

477 Secondly, in Table 2 we also show the results obtained by the multimodal ap-
478 proach when it starts attacking the face segment (randomly initializing the iris
479 section) or iris segment (randomly initializing the face section). As can be ob-
480 served, in both cases the SR is as high as 100% for all the operating points tested.
481 However, the Eff of the attack decreases about 25% when starting with the Sub-
482 Algorithm 2 (Genetic Algorithm) compared to the case of starting with the Sub-
483 Algorithm 1 (Uphill-Simplex). The reason lies on the Eff of the individual Sub-
484 Algorithms. On the left columns of Table 2 (Unimodal Attacks) we can observe
485 that the Eff of the Sub-Algorithm 1 is between 15 and 20 times higher than the
486 Eff of Sub-Algorithm 2 (for a similar number of iterations performed to break an
487 account the number of matchings carried out is significantly higher for the binary
488 attack as was presented in Sects. 3.1 and 3.2). When the multimodal algorithm
489 starts attacking the iris segment, in many occasions it is able to break the system
490 without changing to the face segment. This way, the multimodal attacking algo-
491 rithm can not benefit from the higher Eff of the Sub-Algorithm 1, and has a lower
492 Eff than that achieved when the attack is started against the face section.

493 From the previous observations none of the two main vulnerability scenarios
494 considered for the multimodal attack is clearly better than the other. On the one
495 hand, when it starts attacking the face segment, it is faster but it needs to use both
496 sections of the template to break the system (i.e., face and iris). On the other

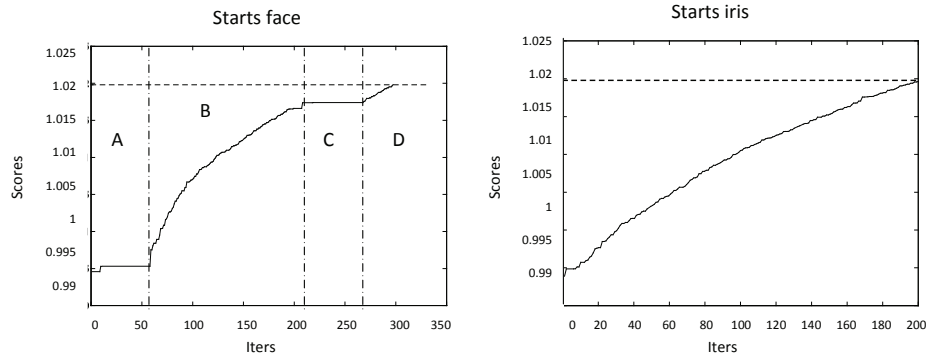


Figure 7: Evolution of the score in each iteration for two broken accounts in the two different scenarios studied: the algorithm starts attacking the face section of the template (left) or the iris section (right). The verification threshold is represented with a dashed horizontal line. In the left plot, the different phases of the algorithm, alternatively attacking the face and iris sections, are marked with letters A-D.

497 hand, when it starts attacking the iris segment, it becomes slower but it has a good
 498 chance of gaining access to the system using just one of the template sections (i.e.,
 499 iris) with the advantage that this may entail in terms of simplification of the attack.

500 In Table 2 we can also observe that the most robust system in terms of Eff and
 501 SR is the unimodal system based on iris and not the multimodal approach as would
 502 be expected. This shows that, as already demonstrated for spoofing attacks [13,
 503 15, 19], although in general multimodal systems offer a better performance than
 504 their unimodal subsystems (for our particular case the EER decreases from 5% to
 505 0.8%), they are not necessarily less vulnerable to software attacks. These results
 506 reinforce the importance of reporting the SR of the attack always in terms of the
 507 operating point at which it was evaluated (i.e., FAR), so that a fair comparison
 508 across different recognition systems may be established.

509 Finally, in Fig. 7 the evolution of the score for each iteration of the algorithm

510 can be observed. On the left, the face section of the template is first attacked, and
511 several areas with different slopes can be observed (marked with letters A, B, C
512 and D), depending on what part of the template is being attacked. In segments
513 A and C, it can also be observed that the algorithm switches to attack the other
514 section of the template after the score remains almost constant for a fixed number
515 of iterations. On the other hand, on the graph on the right, no “steps” can be ob-
516 served on the curve: the attack started attacking the iris section and never changed
517 to the face segment as the template was successfully broken using only the iris
518 part.

519 **7. Countermeasuring the Attack: Score Quantization**

520 Given the high vulnerability of the multimodal system evaluated to the com-
521 bined attacking algorithm proposed, some attack protection needs to be incorpo-
522 rated in order to increase the robustness of the system. When a countermeasure is
523 introduced in a biometric system to reduce the risk of a particular attack, it should
524 be statistically evaluated considering two main parameters:

- 525 • Impact of the countermeasure in the system performance. The inclusion of
526 a particular protection scheme might change the FAR and FRR of a system,
527 and these changes should be evaluated and reported (other performance in-
528 dicators such as speed or computational efficiency might also change, but
529 are not considered here).
- 530 • Performance of the countermeasure, i.e. impact of the countermeasure in
531 the SR and Eff of the attack.

532 It is often argued that a simple account lockout policy (i.e., blocking the user
533 accounts after a number of consecutive unsuccessful access attempts) would be
534 enough to prevent an attack such as the one proposed in the present work. How-
535 ever, such countermeasures still leave the system vulnerable to a spyware-based
536 attack that interlaces its false attempts with the attempts by genuine users (suc-
537 cessful attempts) and collects information over a period of time (i.e. piggyback
538 attack). Furthermore, it may be used by the attacker to perform an account lock-
539 out attack (i.e., the intruder tries to illegally access a great amount of accounts
540 blocking all of them and collapsing the system).

541 In this scenario, a specific design of the matching algorithm can also be im-
542 plemented in order to reduce the effects of this type of threats, providing this way
543 an additional level of security through a biometric-based protection scheme com-
544plementary to other possible non-biometric countermeasures.

545 Among the biometric-based approaches to reduce the effects of hill-climbing
546 attacks, score quantization has been proposed as an effective countermeasure [29].
547 In fact, the BioApi Consortium [45] recommends that biometric algorithms emit
548 only quantized matching scores. Such quantization means that small changes in
549 the randomly generated templates will normally not result in a modification of the
550 matching score, so that the attack does not have the necessary feedback from the
551 system to be carried out successfully.

552 With this precedents, in this section we analyse the performance of score quan-
553 tization as a possible countermeasure against the proposed attack. In the exper-
554 iments we will consider the multimodal system operating at a medium security
555 operating point ($FAR = 0.05\%$). For the combined attack we will assume the
556 same configuration used in the vulnerability assessment experiments.

Table 3: Performance (in terms of SR and Eff) of the combined attack against the system considering different quantization steps (QS), applied before and after the fusion of the scores.

QS		10^{-4}	10^{-3}	10^{-2}	10^{-1}
Before Fusion	SR	100%	100%	0%	0%
	Eff ($\times 10^{-4}$)	1.8932	1.6113	-	-
After Fusion	SR	100%	100%	100%	0%
	Eff ($\times 10^{-4}$)	1.7806	1.7921	1.7470	-

557 Since the global score in this multimodal system is obtained from two previous
558 partial (face and iris) scores that are normalised and then fused, the quantization
559 can take place either before or after this sum or fusion. Both possible schemes are
560 studied in this section.

561 In order to select the appropriate quantization step according to the trade-off
562 that should be met in terms of its impact on the system performance (ideally as
563 small as possible) and on the attack performance (as big as possible), several
564 Quantization Steps (QS) are tested in terms of their corresponding Positive Incre-
565 ment, PI (i.e., percentage of iterations that produced an increase in the similarity
566 score higher than the quantization step considered). The EER of the system with
567 the different QS is computed when the quantization is applied before and after
568 the score fusion. The QS considered range from 10^{-8} and 10^{-1} . For the last QS
569 (10^{-1}), the EER increases considerably (i.e., the QS is too big), while for the re-
570 maining values the performance of the system is not significantly affected. The
571 multimodal attack is therefore repeated applying four QS values, namely: *i*) QS =
572 10^{-4} , *ii*) QS = 10^{-3} , *iii*) QS = 10^{-2} , and *iv*) QS = 10^{-1} . The first three QS values

573 guarantee a similar performance of the system, while the last one can be useful for
574 very high-security applications, when a lower performance of the system might be
575 acceptable if it leads to a much higher protection against the analysed attacks.

576 In Table 3 the results of these experiments are shown. As can be seen, the
577 quantization of the scores is effective as a countermeasure against the combined
578 attacking algorithm presented in this work when it is applied:

- 579 • Before the fusion with a $QS = 10^{-2}$. Since the rounding effect of quantiz-
580 ing the scores and then summing them is bigger than that obtained when
581 fusing the scores before applying the quantization, the performance of the
582 attack decreases more when applying the quantization before the fusion.
583 This leads to a $SR = 0\%$ for the $QS = 10^{-2}$ when the partial scores are
584 quantized before fusing them.
- 585 • Before or after the fusion with a $QS = 10^{-1}$. With this QS , the system is able
586 to stop the attack regardless of the point where the scores are quantized. As
587 in the previous case, the attack does not receive the necessary feedback from
588 the system on whether it has managed to increase or not the similarity score,
589 and thus fails to achieve its objective.

590 In both cases listed above, no account is broken, while for the remaining trials
591 the SR of the attack is still 100%, only decreasing its Eff (i.e., more comparisons
592 are needed to break an account). However, while the performance of the system is
593 not considerably affected in the first case ($EER = 1.37\%$), it is barely acceptable
594 with a $QS = 10^{-1}$: the EER is as high as 32.06%.

595 8. Conclusions

596 In this work, we have presented and evaluated the first software attack against
597 multimodal biometric systems. As case study, we have tested it on a system based
598 on face and iris, a trait combination regarded as user-friendly: the features of both
599 traits may be extracted from images that can be captured at the same time, being
600 the acquisition process transparent to the user. The attacking algorithm shows a
601 remarkable performance, thus proving the vulnerabilities of multimodal systems
602 to this type of attacks. Furthermore, the multimodal system has not presented
603 an improvement in the security level against this kind of attack compared to the
604 face and iris modules on their own. This fact confirms what previous studies
605 on spoofing attacks pointed out: even though multimodal systems recognition
606 performance is higher, they do not necessarily increase the robustness of unimodal
607 approaches to external attacks.

608 The quantization of the scores given by the matcher is analysed as a possible
609 countermeasure. Two different approaches are studied and compared: the partial
610 scores can be quantized before fusing them, or the final score can be quantized
611 after the fusion. The first scenario leads to a null success rate without affecting
612 the verification performance of the system, being thus a suitable countermeasure
613 for the proposed attack. The second case also protects the system against the
614 attack but at the cost of drastically reducing its verification performance.

615 Research works such as the one presented in this article pretend to bring some
616 insight into the difficult problem of biometric security evaluation through the sys-
617 tematic study of biometric systems vulnerabilities and the analysis of effective
618 countermeasures that can minimize the effects of the detected threats, in order to
619 increase the confidence of the final users in this rapidly emerging technology.

620 References

- 621 [1] S. Oviatt, P. Cohen, Multimodal interfaces that process what comes natu-
622 rally, *Commun. ACM* 43 (2000) 45–53.
- 623 [2] S. Oviatt, Ten myths of multimodal interaction, *Commun. ACM* 42 (1999)
624 74–81.
- 625 [3] A. K. Jain, A. Ross, S. Pankanti, Biometrics: a tool for information security,
626 *IEEE Trans. on Information Forensics and Security* 1 (2006) 125–143.
- 627 [4] J. Wayman, A. Jain, D. Maltoni, D. Maio, *Biometric systems. Technology,*
628 *design and performance evaluation*, Springer, 2005.
- 629 [5] B. Schneier, Inside risks: the uses and abuses of biometrics, *Commun. ACM*
630 42 (1999) 136.
- 631 [6] J. Galbally, J. Fierrez, F. Alonso-Fernandez, M. Martinez-Diaz, Evaluation
632 of direct attacks to fingerprint verification systems, *Telecommunication Sys-*
633 *tems, Special Issue on Biometrics* 47 (2011) 243–254.
- 634 [7] T. Matsumoto, Artificial irises: importance of vulnerability analysis, in:
635 *Proc. 2nd Asian Biometrics Workshop*, 2004.
- 636 [8] J. Galbally, C. McCool, J. Fierrez, S. Marcel, On the vulnerability of face
637 verification systems to hill-climbing attacks, *Pattern Recognition* 43 (2010)
638 1027–1038.
- 639 [9] U. Uludag, A. Jain, Attacks on biometric systems: a case study in fin-
640 gerprints, in: *Proc. SPIE Seganography and Watermarking of Multimedia*
641 *Contents VI*, 2004, volume 5306, pp. 622–633.

- 642 [10] Tabula Rasa, Trusted biometrics under spoofing attacks (tabula rasa), 2010.
- 643 [11] B. Schneier, Secrets and lies, Wiley, 2000.
- 644 [12] A. Kerckhoffs, La cryptographie militaire, Journal des
645 Sciences Militaires 9 (1883) 5–83. Available on-line at
646 <http://www.petitcolas.net/fabien/kerckhoffs>.
- 647 [13] Z. Akhtar, S. Kale, N. Alfarid, Spoof attacks in multimodal biometric sys-
648 tems, in: Proc. International Conference on Information and Network Tech-
649 nology (IPCSIT), 2011, volume 4, IACSIT Press, 2011, pp. 46–51.
- 650 [14] G. Chetty, M. Wagner, Audio-visual multimodal fusion for biometric person
651 authentication and liveness verification, in: Proc. NICTA-HCSNet Multi-
652 modal User Interaction Workshop (MMUI), 2005.
- 653 [15] R. N. Rodrigues, L. L. Ling, V. Govindaraju, Robustness of multimodal
654 biometric fusion methods against spoof attacks, Journal of Visual Languages
655 and Computing 20 (2009) 169–179.
- 656 [16] R. Rodrigues, N. Kamat, V. Govindaraju, Evaluation of biometric spoof-
657 ing in a multimodal system, in: Proc. IEEE International Conference on
658 Biometrics: Theory Applications and Systems (BTAS), 2010.
- 659 [17] Z. Akhtar, N. Alfarid, Secure learning algorithm for multimodal biometric
660 systems against spoof attacks, in: Proc. International Conference on Infor-
661 mation and Network Technology (IPCSIT), 2011, volume 4, IACSIT Press,
662 2011, pp. 52–57.

- 663 [18] J. Hämmerle-Uhl, K. Raab, A. Uhl, Attack against robust watermarking-
664 based multimodal biometric recognition systems, in: Proc. of the COST
665 2101 European conference on Biometrics and ID management (BioID),
666 2011, LNCS-6583, 2011, pp. 25–36.
- 667 [19] P. Johnson, B. Tan, S. Schuckers, Multimodal fusion vulnerability to non-
668 zero effort (spoof) attacks, in: Proc. Workshop on Information Forensics and
669 Security (WIFS), 2010.
- 670 [20] E. Marasco, Secure Biometric Systems, Ph.D. thesis, University of Naples
671 Federico II, 2010.
- 672 [21] M. Martinez-Diaz, J. Fierrez, J. Galbally, J. Ortega-Garcia, An evaluation
673 of indirect attacks and countermeasures in fingerprint verification systems,
674 Pattern Recognition Letters 32 (2011) 1643–1651.
- 675 [22] Y. Wang, T. Tan, A. K. Jain, Combining face and iris biometrics for iden-
676 tity verification, in: Proc. of Int. Conf. on Audio- and Video-Based Person
677 Authentication (AVBPA), 2003, pp. 805 – 813.
- 678 [23] X. Zhang, Z. Sun, T. Tan, Hierarchical fusion of face and iris for personal
679 identification, in: Proc. International Conference on Pattern Recognition
680 (ICPR), 2010.
- 681 [24] J.-Y. Gan, J.-F. Liu, Fusion and recognition of face and iris feature based
682 on wavelet feature and kfda, in: Proc. of the International Conference on
683 Wavelet Analysis and Pattern Recognition (ICWAPR), 2009.
- 684 [25] J.-Y. Gan, Y. Liang, A method for face and iris feature fusion in identity

- 685 authentication, International Journal of Computer Science and Network Se-
686 curity (IJCSNS) 2 (2006) 135–138.
- 687 [26] S. Oviatt, R. Coulston, S. Tomko, B. Xiao, R. Lunsford, M. Wesson,
688 L. Carmichael, Toward a theory of organized multimodal integration pat-
689 terns during human-computer interaction, in: Proc. Int. Conf. on Multimodal
690 Interaction, 2003.
- 691 [27] N. Ratha, J. H. Connell, R. M. Bolle, An analysis of minutiae matching
692 strength, in: Proc. IAPR on Audio- and Video-Based Person Authentication
693 (AVBPA), 2001, Springer LNCS-2091, 2001, pp. 223–228.
- 694 [28] B. Tan, Assessing and reducing spoofing vulnerability for multimodal and
695 fingerprint biometrics, Ph.D. thesis, Clarkson University, 2009.
- 696 [29] A. Adler, Images can be regenerated from quantized biometric match score
697 data, in: Proc. Canadian Conference on Electrical and Computer Engineer-
698 ing (CCECE), 2004, pp. 469–472.
- 699 [30] J. A. Nelder, R. Mead, A simplex method for function minimization, Com-
700 puter Journal 7 (1965) 368 – 313.
- 701 [31] M. Gomez-Barrero, J. Galbally, J. Fierrez, J. Ortega-Garcia, Hill-climbing
702 attack based on the uphill simplex algorithm and its application to signa-
703 ture verification, in: Proc. European Workshop on Biometrics and Identity
704 Management (BioID), 2011, LNCS-6583, 2011, pp. 83–94.
- 705 [32] A. Brindle, Genetic Algorithms for Function Optimization, Ph.D. thesis,
706 University of Alberta, Edmonton, 1981.

- 707 [33] J. E. Baker, Reducing bias and inefficiency in the selection algorithm, in:
708 Proc. International Conference on Genetic Algorithms and their Application
709 (ICGAA), 1987, L. Erlbaum Associates Inc., 1987, pp. 14 – 21.
- 710 [34] M. Gomez-Barrero, J. Galbally, P. Tome-Gonzalez, J. Fierrez, On the vulner-
711 ability of iris-based systems to software attacks based on genetic algorithms,
712 in: Proc. Iberoamerican Conf. on Pattern Recognition (CIARP), 2012.
- 713 [35] L. Masek, P. Kovesi, MATLAB Source Code for a Biometric Identification
714 System Based on Iris Patterns, Master’s thesis, School of Computer Science
715 and Software Engineering, University of Western Australia, 2003.
- 716 [36] M. A. Turk, A. P. Pentland, Face recognition using eigenfaces, in: Proc.
717 IEEE Conference on Computer Vision and Pattern Recognition (CVPR),
718 1991, pp. 586–591.
- 719 [37] J. Phillips, P. Flynn, et al., Overview of the face recognition grand challenge,
720 in: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR),
721 2005, pp. 947–954.
- 722 [38] V. Ruiz-Albacete, P. Tome-Gonzalez, et al., Direct attacks using fake images
723 in iris verification, in: Proc. European Workshop on Biometrics and Identity
724 Management (BioID), 2008, LNCS-5372, 2008, pp. 181–190.
- 725 [39] J. Daugman, How iris recognition works, Circuits and Systems for Video
726 Technology, IEEE Transactions on 14 (2004) 21–30.
- 727 [40] A. K. Jain, K. Nandakumar, A. Ross, Score normalization in multimodal
728 biometric systems, Pattern Recognition 38 (2005) 2270–2285.

- 729 [41] A. Ross, A. K. Jain, Information fusion in biometrics, *Pattern Recognition*
730 *Letters* 24 (2003) 2115 – 2125.
- 731 [42] J.Ortega-Garcia, J.Fierrez, F.Alonso-Fernandez, J.Galbally, M.R.Freire,
732 J.Gonzalez-Rodriguez, C.Garcia-Mateo, J.-L.Alba-Castro, E.Gonzalez-
733 Agulla, E.Otero-Muras, S.Garcia-Salicetti, L.Allano, B.Ly-Van, B.Dorizzi,
734 J.Kittler, T.Bourlai, N.Poh, F.Deravi, M.W.R.Ng, M.Fairhurst, J.Hennebert,
735 A.Humm, M.Tistarelli, L.Brodo, J.Richiardi, A.Drygajlo, H.Ganster,
736 F.M.Sukno, S.-K.Pavani, A.Frangi, L.Akarun, A.Savran, The multi-scenario
737 multi-environment BioSecure multimodal database (BMDB), *IEEE Trans.*
738 *on Pattern Analysis and Machine Intelligence* 32 (2010) 1097–1111.
- 739 [43] A. Mayoue, B. Dorizzi, L. Allano, G. Chollet, J. Hennebert, D. Petrovska-
740 Delacretaz, F. Verdet, Guide to biometric reference systems and performance
741 evaluation, Springer, pp. 327–372. (2009).
- 742 [44] ANSI, 2001. ANSI X9.84-2001, Biometric Information Management and
743 Security.
- 744 [45] BioAPI, The BioAPI consortium, 2009. <http://www.bioapi.org>.