Iris Based Identity Verification Robust to Sample Presentation Security Attacks

Inmaculada Tomeo-Reyes, Vinod Chandran

School of Electrical Engineering and Computer Science, Queensland University of Technology, Brisbane, Queensland, Australia

Received: 10 April 2013; Accepted: 15 May 2013

Abstract

Iris based identity verification is highly reliable but it can also be subject to attacks. Pupil dilation or constriction stimulated by the application of drugs are examples of sample presentation security attacks which can lead to higher false rejection rates. Suspects on a watch list can potentially circumvent the iris based system using such methods. This paper investigates a new approach using multiple parts of the iris (instances) and multiple iris samples in a sequential decision fusion framework that can yield robust performance. Results are presented and compared with the standard full iris based approach for a number of iris degradations. An advantage of the proposed fusion scheme is that the trade-off between detection errors can be controlled by setting parameters such as the number of instances and the number of samples used in the system. The system can then be operated to match security threat levels. It is shown that for optimal values of these parameters, the fused system also has a lower total error rate.

Keywords: Iris; Verification; Parts; Samples; DET

1. Introduction

The iris is a highly accurate biometric trait and iris recognition systems have been proven to be highly reliable [1]. Detection errors such as false acceptances and false rejections are very low for these systems; however, they can be subject to security attacks that increase them. By increasing the false rejection rate, for example, a suspect can prevent being matched to a watch list and bypass an iris based system. Sample presentation attacks are common approaches used to degrade the performance of iris recognition in this manner.

In general, when biometric systems are in operation and the threat level changes, controlling the trade-off between detection error rates can reduce the impact of an attack. Traditional unimodal biometric systems do not allow achieving this, since just a single threshold can be varied. In multimodal or fused systems, on the contrary, this can be achieved through systematic adjustment of a set of parameters. Recent investigations [2, 3] have addressed these issues using text-dependent speech. Results show that in a sequential decision fusion architecture with multiple classifiers (instances) and multiple attempts (samples) it is possible to control the trade-off between false accept and false reject errors at each classifier using the number of instances and the number of samples considered in the architecture. The fact that decision level fusion [4, 5] is considered in the above-mentioned work, has the advantage that single modality based systems can combine quite easily even if these systems have been developed independently, possibly by different companies, and internal workings are not made public, something very common in the case of many commercial off-the-shelf (COTS) biometric matchers. There is, however, some trade-off through loss of information in the threshold operation that makes each individual decision but this is not critical if individual systems working with single modalities have been designed to make near-optimal use of all information available to them. These considerations are also valid in the case of multiple classifiers designed using the same modality but different instances containing potentially different information. For example, voice recognition systems can be text-dependent and work with different text (of the utterance) as different instances. Iris recognition systems may work with the left and right iris or with parts of a given iris as each instance.

This paper is aimed at investigating the viability of adapting and applying a multi-instance, multi-sample sequential decision fusion architecture to iris recognition in the context of preventing sample presentation attacks. How sample presentation attacks, and more specifically, obfuscation attacks, affect iris recognition systems is explained in section 2. Details of the fusion architecture can be found in Section 3. Section 4 describes the methodology used to evaluate the fused system performance and the results.

2. Security Threats to Biometric Systems: Sample Presentation Attacks

When working with biometric systems, it is very important to consider the potential security threats, as they can lead to security failures. Security failures can occur from intrinsic limitations of the system, or due to explicit attacks. Any biometric system, regardless of the trait, comprises at least four different subsystems: data acquisition, pre-processing, feature extraction and comparison [6]. Each subsystem has different points of attack, with one or more potential methods of exploitation of each. Some of the early work by Ratha et al. [7] identified eight possible points of attack. Further work by Jain et al. [8] sought to refine this approach. Wayman [9] added the storage block to the general biometric schema, allowing a more detailed analysis of the different points of attack. Combining elements of previous works, Bartlow and Cukic [10] extended this research by adding three new components: administrative supervision, information technology environment and token presentation. Common Criteria [11, 12], an international standard used for computer security, particularly by governments, also defines an extended general biometric schema that includes storage and three different administrative subsystems, apart from the four basic subsystems. Figure 1 shows a biometric system at the subsystem level and highlights the potential attack points according to Common Criteria. The general threats that need to be considered when evaluating biometric systems for vulnerabilities are shown in the adjacent table. The outline numbers correspond to locations identified by the numbers in Figure 1. A complete description of the specific kind of attacks corresponding to each threat can be found in [12]. Among all possible threats shown in Figure 1, this research work will focus on user/capture threats (type 2), also known as attacks at the user interface level or sample presentation attacks.

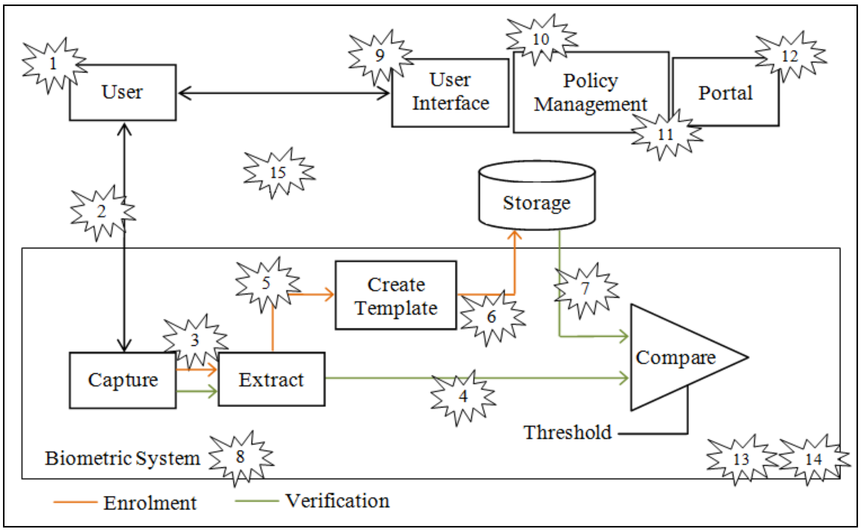

1	User Threats
2	User/Capture Threats (Sample presentation attacks)
3	Capture/Extraction Threats
4	Extraction/Comparison Threats during Verification
5	Extraction/Template Storage Threats during Enrolment
6	Template Storage Threats
7	Template Retrieval Threats
8	Administrator/Resource Manager Threats
9	User/Policy Management Threats
10	Policy Management Threats
11	Threats to Policy Management/Portal
12	Portal Threats
13	Threats to all hardware components
14	Threats to all software or firmware components
15	Threats to all connections (including network threats)

Figure 1. General threats for biometric systems

2.1 Sample presentation attacks: overview

In general, any attempt by an attacker to break into the system by presenting an altered biometric sample can be considered a sample presentation attack. These attacks can be categorized into three main groups [6]:

(a) *Impersonation.* In impersonation attacks impostors pose themselves as an authorized user in an attempt to intrude the system. If the impersonation is casual, the identity to attack is randomly chosen and the impostor's biometric characteristics are not modified. If the impersonation is targeted, a specific identity is attacked. In this case, the impostor may modify his own behaviour (e.g. voice, signature or gait) or physiology (e.g. face or hand) in an attempt to match the identity under attack. Targeting an identity which is known to be easier to impersonate (e.g. weak biometric template) or whose biometric characteristics are known to be similar to the impostor's (e.g. twin) are other options.

(b) *Spoofing.* Spoofing attacks involve the presentation of a counterfeit biometric that does not come from a live person. The usage of fake or artificial traits (e.g. gummy finger, recorded voice, etc.) or non-live samples (e.g. dismembered finger) from legitimate users are typical spoofing attacks [13, 14].

(c) *Obfuscation or disguise.* Obfuscation attacks occur when the attacker deliberately changes his biometric characteristic in order to avoid being recognized by the system. Intentionally presenting a noisy, poor-quality or null biometric sample that may not match the template, using disguises or plastic surgery in the case of face, or applying techniques to obliterate fingerprints (e.g. abrasion, cutting or burning) are common examples of obfuscation.

The main problem for an impostor to successfully perform impersonation or spoofing attacks is that it is necessary to have a good copy or some prior knowledge of the biometric trait corresponding to the identity to be attacked. The difficulty of this depends strongly on the trait. Obfuscation attacks, on the contrary, can be easily carried out regardless of the trait and no previous knowledge is generally required. Apart from that, these attacks can be used to illegitimately gain access to a system by circumventing the main and most secure subsystem and taking advantage of less secure fallback mechanisms if they exist. In such case they can be considered as dangerous as impersonation or spoofing. According to this, the necessity of minimizing the impact of obfuscation attacks is clear and so, we will focus on these attacks in this paper.

2.2 Obfuscation attacks and iris recognition

Iris patterns are unique regardless of any genetic relationship between their possessors. Even the two iris patterns of the same individual are unique and structurally different. The human iris is also stable and not changeable, except in cases of injury. Because of these two important characteristics, uniqueness and stability, iris recognition is considered a highly reliable method of identification and/or verification. Iris recognition is in fact one of the most accurate biometric techniques [1], but in spite of that, obfuscation attacks degrade its performance [15]. Threats associated with obfuscation attacks in iris recognition can be grouped into three categories:

(a) *Intentional presentation of a noisy or poor-quality sample.* The first possibility involves intentionally presenting a noisy or poor-quality biometric sample that may not match the attacker template in the database. Blinking, deviating the gaze and using glasses are typical ways to achieve this. It is important to note that not all cases in which noisy or poor quality iris samples are captured are potential disguise attacks. Both noise and poor image quality are common problems in iris recognition systems and a lot of research work has already been done to improve the performance of these systems under such conditions [16]. However, unlike the architecture proposed in this research work, methods developed for this purpose are not aimed to prevent obfuscation attacks.

(b) *Artificially provoked iris alterations*. Given the anatomic characteristics of the iris, altering it in any way is extremely difficult. However, two types of iris alterations can be easily provoked: mydriasis and miosis. Mydriasis is an excessive dilation of the pupil arising from disease, trauma or the use of drugs or alcohol. It can also be artificially provoked by using a mydriatic agent in the form of eyedrops. Non-elastic deformations of the iris occur as the pupil dilates. Miosis, on the other hand, is an excessive constriction of the pupil which can also be artificially provoked by using a miotic agent. The effect of pupil dilation on iris recognition has been addressed before [1, 17]; however, this research did not extend to extreme cases of pupil dilation or constriction (either pathological or artificially provoked) or changes in the threat level.

(c) *Occlusion of the iris.* The iris can be easily occluded by using cosmetic lenses, which are contact lenses with a pattern printed or painted on them. Total occlusion can be achieved by using colour, fantasy or prosthetic lenses. As they are opaque, the corresponding iris texture becomes unavailable. Partial occlusion is also possible by using lenses in which only a black inner circle simulating the pupil has been printed or painted.

In order to demonstrate how these attacks affect iris recognition performance, an iris recognition algorithm inspired by Daugman's works [18] has been used here (see subsection 4.2). Images from the three previous categories have been used to calculate the detection error rates corresponding to the system threshold under normal operating conditions (see Figure 2). In each case, the corresponding detection error trade-off (DET) curve has been added to illustrate the fact that reducing the impact of the attack is not possible by just varying a single threshold of the base system because it results in a trade-off. A multi-part, multi-sample architecture can be used for that purpose instead. Results in Figure 2 clearly illustrate the false rejections increment typical from obfuscation attacks.

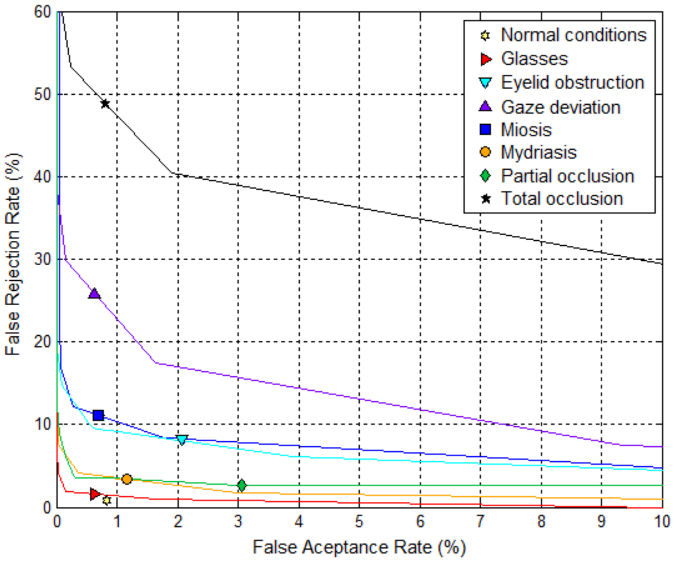

Figure 2. Detection error rates for different obfuscation attacks considering system threshold under normal operating conditions

3. Multi-biometric Fusion Approach

A multi-instance, multi-sample sequential decision fusion architecture recently developed by Nallagatla and Chandran [2, 3], allows controlled trade-off between false accept and false reject errors. This architecture is particularly suited to users of commercial systems who may have a number of different systems at their disposal but access only to decisions rather than scores, features or internal algorithms. The architecture has so far been tested with short speech utterances. In this paper, it is further developed and applied to an iris recognition system, with special emphasis on preventing sample presentation attacks.

3.1 Multi-part fusion scheme

In the context of iris recognition, instance usually refers to each of the irises of an individual. In this research work, this concept is generalized, and different parts of the same iris image are considered as the different instances. Three main reasons motivate this decision:

(a) Deliberate iris changes provoked to perform obfuscation attacks are applicable to both irises in most cases, so using them as the two only instances does not seem to be useful.

(b) Since the number of instances is one of the key factors to control detection errors, the possibility of considering more than just two inputs (left and right irises) makes the system more versatile and errors easier to control.

(c) Although a part based approach is mainly used here to make the base architecture robust to obfuscation attacks, other advantages of iris recognition using parts such as dealing with noise [19], cancellable iris biometrics [20] or efficiency improvement [21] also apply.

There are many schemes for dividing an iris into parts. Concentric rings of equal width, as shown in Figure 3, have been chosen in this paper. Iris templates are built for each ring and classifiers for the different rings form the ‘*n*’ stages of the sequential decision fusion architecture. In this multi-part fusion scheme, the user’s claim is only accepted if the decisions from all stages are *accept*. This is equivalent to applying an AND decision fusion rule to determine the acceptance of a claim. This fusion method effectively reduces the false acceptances; however, it can increase the number of false rejections (see Table 1).

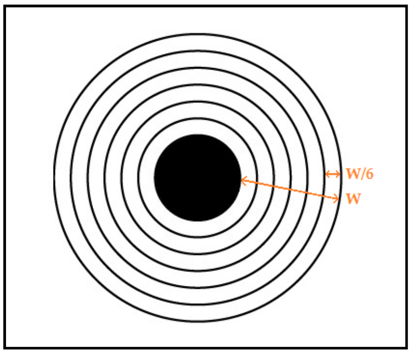

Figure 3. Multi-part fusion: concentric rings used as parts

3.2 Multi-sample fusion scheme

In order to reduce false rejections, multiple iris samples can be used. These samples are divided into rings in the same manner. In the multi-sample fusion scheme, up to '*m*' samples can be presented at any given stage (part classifier) if the decision is *reject*. If the maximum number of allowable samples ('*m*') is exceeded, the claim is rejected. Acceptance of a claim at any given stage is equivalent to an OR decision fusion rule among samples. This fusion method helps in reducing false rejections, but it can increase false acceptances since the impostor is given additional chances for verification (see Table 1).

3.1 Multi-part, multi-sample fusion scheme: proposed architecture

The proposed architecture, resulting from combining the multi-part and multi-sample fusion schemes, can be seen in Figure 4. It consists of '*n*' classifiers arranged sequentially and allows up to '*m*' samples presentation (dashed line) in case any of the user's claims is rejected by the system (d=0).

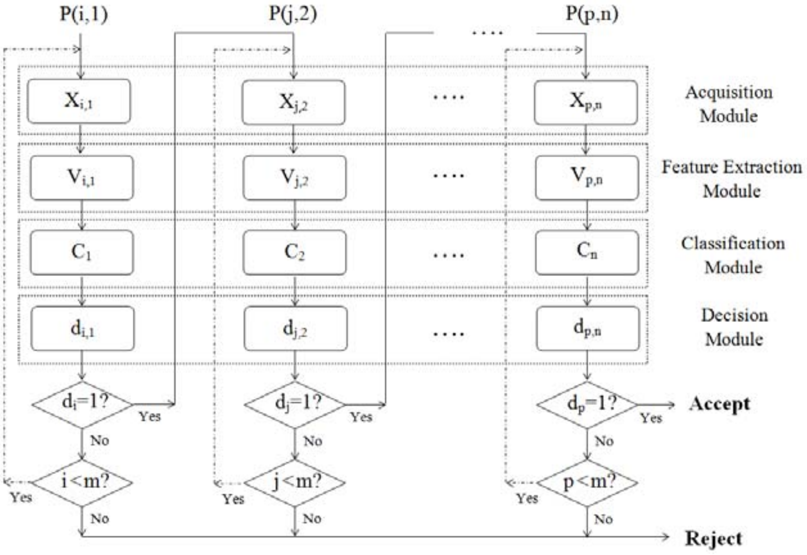

Figure 4. Multi-part, multi-sample decision fusion scheme

False acceptance rate (FAR) and false rejection rate (FRR) of each of the fused systems in terms of false acceptance and false rejection for each independent attempt (α and ρ) when considering '*n*' stages and '*m*' sample presentation attempts can be calculated according to the equations shown in Table 1 [2].

Table 1. Detection error rates of fused schemes: multi-part, multi-sample and both

Fusion scheme	Detection Error Rates: FAR (α) and FRR (ρ)
Multi-part	$\alpha_{fused} = \alpha^n \quad (1)$ $\rho_{fused} = \rho + (1-\rho)\rho + (1-\rho)^2\rho + \cdots + (1-\rho)^{n-1}\rho \quad (2)$
Multi-sample	$\alpha_{fused} = m\alpha \quad (3)$ $\rho_{fused} = \rho^m \quad (4)$
Multi-part and multi-sample	$\alpha_n^m = (m\alpha)^n \quad (5)$ $\rho_n^m = \rho^m + (1-\rho^m)\rho^m + \cdots + (1-\rho^m)^{n-1}\rho^m \quad (6)$

From the above equations it is clear that when considering multi-part fusion, since α is less than 1, the resultant FAR (α_{fused}) decreases multiplicatively with the number 'n' of classifiers or parts used. The fused FRR, however, increases as the addition of the term $(1-\rho)^{n-1}\rho$ from 1 to 'n' – in fact equation (2) can be reduced to $\rho_{\text{fused}} \approx n\rho$ when $\rho \ll 1$ but such simplification is not appropriate when dealing with obfuscation attacks, since false rejections are significantly increased with respect to normal operating conditions. According to this, since multiplicative changes are faster than additive ones, the reduction in the FAR is faster than the increase in the FRR. Just the opposite occurs in the case of multi-sample fusion – the resultant FRR decreases multiplicatively with the number 'm' of attempts, while the FAR increases additively. When combining both behaviours, by choosing appropriate values of 'n' and 'm', a desired trade-off between decision error rates can be achieved that reduces the impact of obfuscation attacks. It should be noted that error rates represented by α and ρ in these equations are between 0 and 1; however, percentages will be used from Section 4 on.

4. Tests and Results

4.1 Database

A total of 1476 iris images, acquired at a resolution of 640x468 pixels, constitute the iris dataset used to perform the tests. Images were taken from 59 different participants with ages between 16 and 70 years old. The sensor used for the data collection is the IG-AD100, a dual eye auto-focus camera which works in the near infrared (NIR) wavelength. It has native built-in passive, behavioural and dynamic countermeasures for eye liveness detection, but they were deactivated with the aim of avoiding any restriction when capturing the data.

For generating noisy and poor quality samples, images were obtained with the effects of wearing glasses, blinking and deviating the gaze to any direction (up, down, left and right). For artificially provoked iris alterations, a miotic/mydriatic agent in the form of eyedrops was instilled to participants. Three different types of cosmetic lenses (colour, fantasy and prosthetic) were used to completely occlude the iris. In the case of partial occlusion, images were synthetically created from the normal conditions image subset by increasing the pupil radius from R_p to $1.75 \cdot R_p$ (this value can be adjusted for more/less severe degradation). However, the same effect could have been achieved by using lenses in which only a black inner circle simulating the pupil has been printed or painted. Examples of all these cases can be seen in Figure 5.

Although the initial dataset size was larger than the 1476 images initially stated, around 30% of the images were discarded because of segmentation errors. The increase in segmentation errors is an important effect of sample presentation attacks that also degrades recognition performance, but these images are avoided here because the robustness of the architecture to iris texture changes as a result of these attacks is being evaluated, distinct from segmentation errors. In order to validate the results, a train set was first created to find the classifiers thresholds. Results were later obtained using unseen images from the test set (see Table 2).

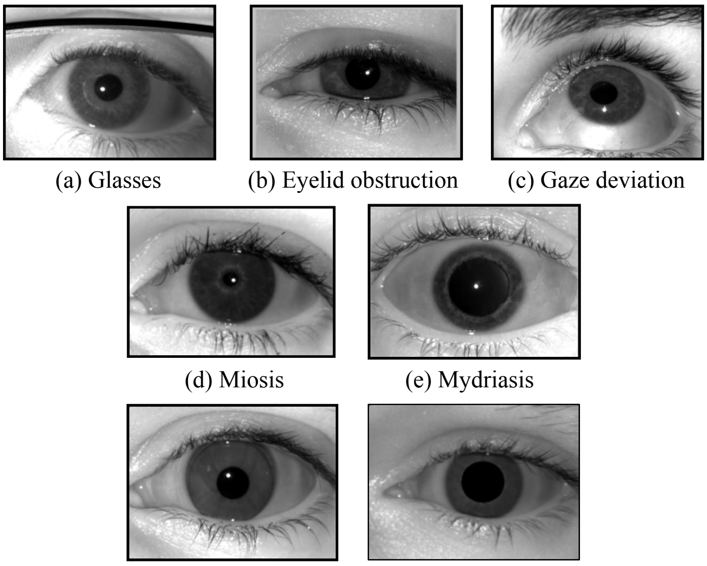

(a) Glasses (b) Eyelid obstruction (c) Gaze deviation

(d) Miosis (e) Mydriasis

(f) Prosthetic lens (total occlusion) (g) Simulated partial occlusion

Figure 5. Sample images: noise or poor quality (a-c), artificial alterations (d-e) and occlusion (f-g)

Table 2. Iris dataset

		Users	Images	Train set	Test set
Normal conditions		59	177	59	118
Noise or poor quality	Glasses	37	215	37 · 3 = 111	104
	Eyelid obstruction	15	114	15 · 3 = 45	69
	Gaze deviation	15	161	15 · 3 = 45	116
Artificial alterations	Mydriasis	14	229	14 · 3 = 42	187
	Miosis	10	248	10 · 3 = 30	218
Occlusion	Partial	59	177	59	118
	Total	9	155	9 · 3 = 27	128

4.2 Iris recognition base algorithm

The implemented iris recognition base algorithm is based on Daugman's approach [18], although some modifications have been introduced to improve computational efficiency. Regarding the image pre-processing stage, the black hole search method [22] is first used to locate the pupil (course search). A simplified version of Daugman's integro-differential operator is then used (fine search) for defining the contours. In order to detect and mask eyelids to eliminate the non-biometric information from the image, a combination of Canny´s edge detection algorithm [23] and Deriche et al. algorithm [24] is used. In the case of the feature extraction stage, convolution with 2-dimensional Gabor filters is considered to extract the texture from the normalized iris image. Filtering is performed by sectors (normalized image is divided into 12 rings and 256 sectors per ring) and only the usable iris area is considered for normalization. Binary coding of the coefficients obtained after filtering and template matching, based on Hamming distance, are carried out as exposed by Daugman. Using an iris dataset composed of 177 images captured under normal conditions from 59 users, the Equal Error Rate (EER) value obtained for a Hamming distance threshold of 44.53 is 0.87%.

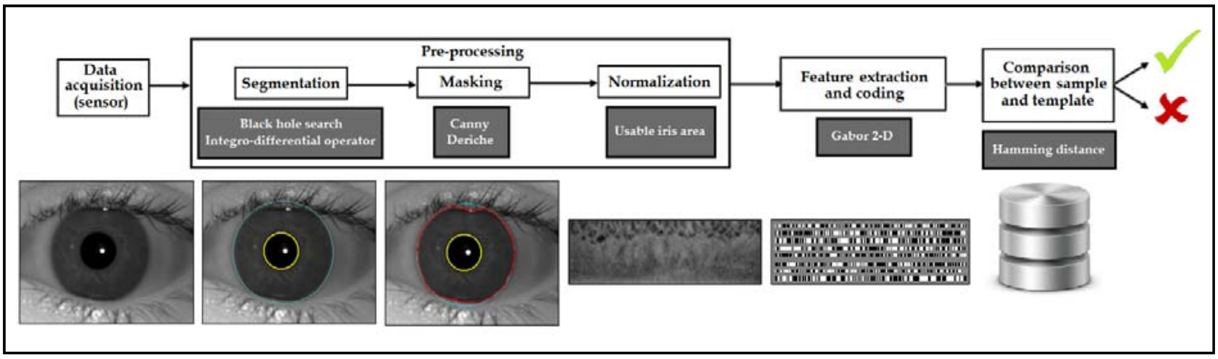
Data acquisition (sensor)
Pre-processing
Segmentation
Masking
Normalization
Black hole search
Integro-differential operator
Canny
Deriche
Usable iris area
Feature extraction and coding
Gabor 2-D
Comparison between sample and template
Hamming distance

Figure 6. Iris recognition base algorithm

4.3 Results

To evaluate the effectiveness of the proposed multi-part, multi-sample fusion scheme and justify the final result, each individual scheme is first evaluated. The EER from the train subset when using the entire iris is used as a reference. The train subset is also used to fix the threshold of each individual classifier, this being the EER-based threshold. Since the threshold of each classifier is fixed in advance, DET curves cannot be calculated. Instead, different values of detection error rates are obtained by progressively increasing the number of parts/classifiers (multi-part) or samples (multi-sample). These values are then connected together to indicate which values belong to the same experiment.

Results for *multi-part fusion* are obtained by progressively increasing the number '*n*' of rings or parts. Four different values of '*n*' have been considered in all cases (n = 2, 3, 4 and 6), being the width of each ring equal to the width of the whole normalized iris divided by '*n*' (see Figure 3). Rings are always processed in the same order, inner to outer. The value of '*n*' defines the maximum number of classifiers in the sequential chain. In the case of *multi-sample fusion*, a maximum of three sample presentations (m = 3) are allowed to verify the user in all cases but in the case of partial occlusion, in which only two sample presentations (m = 2) are possible since there are no more samples available in the database. Samples are chosen randomly from the test data set, and by repeatedly and randomly selecting sets of *'m'* samples, average error rates and standard deviations are estimated and presented in Tables 3-5. To be consistent with the multi-instance experiment, the only classifier existing in this scheme (n = 1) corresponds to the inner ring of the iris (the nearest to the pupil). Finally, results obtained when *integrating the multi-part and multi-sample schemes* are calculated following the previous requirements. For clarity purposes, just the cases in which halves and thirds of the iris (n = 2, 3) are used as parts are considered in Figures 7c-9c, although all cases (n = 2, 3, 4 and 6) have been taken into account in Tables 3-5. Results for all three categories of obfuscation attacks are presented and analysed next.

A. Intentional presentation of a noisy or poor-quality sample

In all three cases within this category (glasses, eyelid obstruction and gaze deviation), results when considering multi-part fusion show the same behaviour: while FAR decreases with the number of parts, FRR increases (see Figure 7 a.1-a.3). According to equations (1) and (2), since the FAR decreases multiplicatively, its reduction is faster than the increase in the FRR. The fast decreasing of the FAR is the main reason why no more than 3 stages are usually required to reduce false acceptances to values very close to 0. An analogous behaviour can be observed in Figure 7 b.1-b.3 for an increasing number '*m*' of samples (m = 1, 2 and 3) in the multi-sample fusion scheme. In this case, the more samples considered the lower value of the FRR, being the FRR reduction faster than the FAR increasing. It may also be noted that a maximum of 3 samples is enough to reduce the FRR values to almost 0.

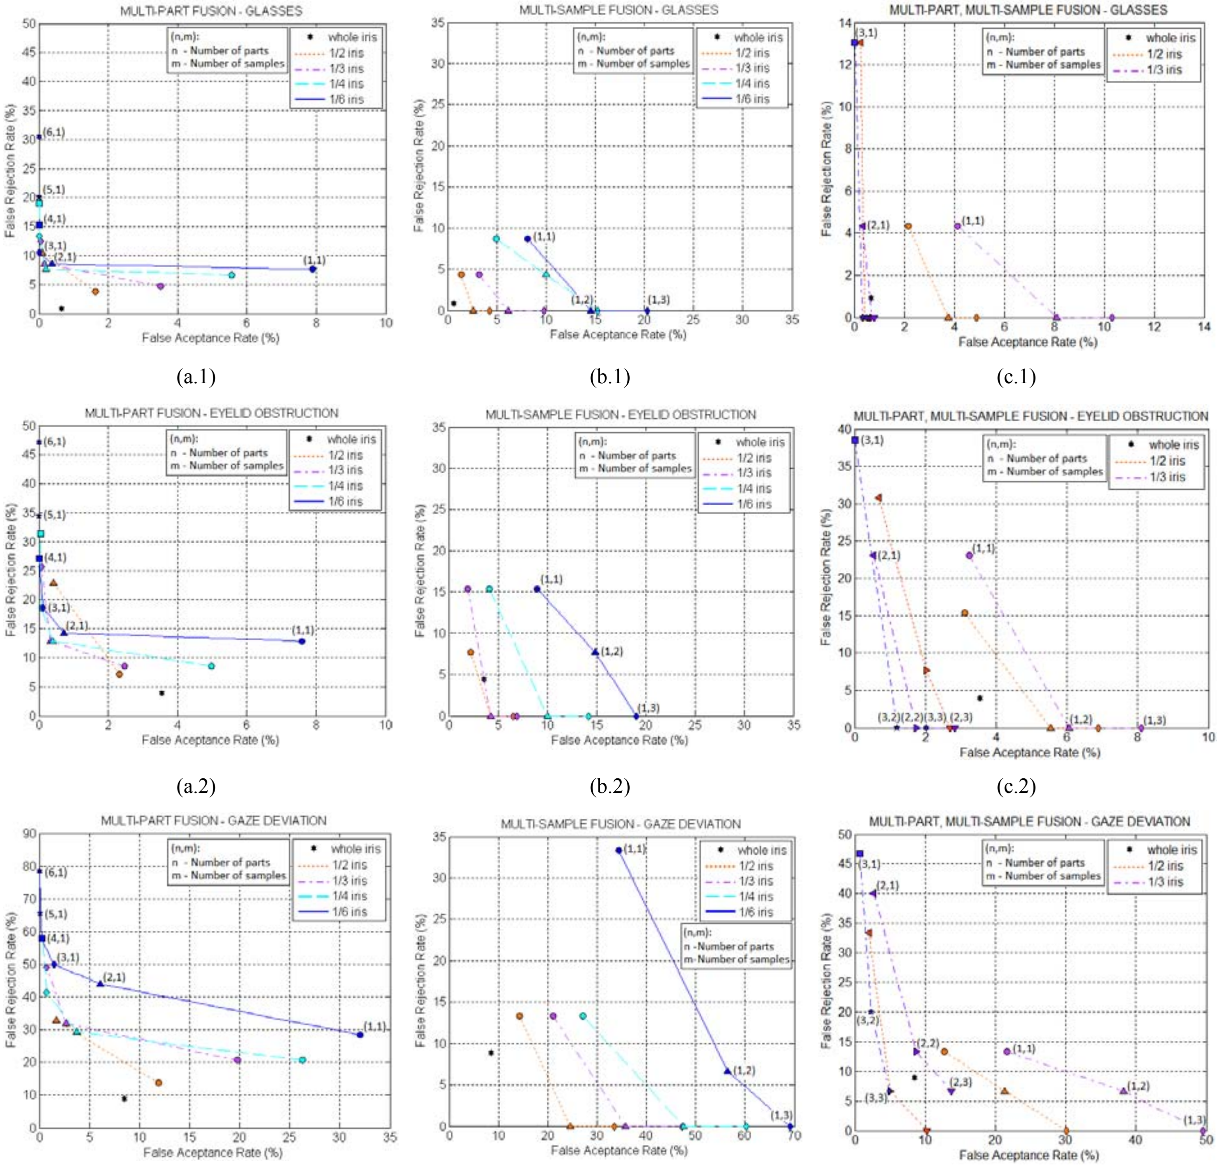

(a.1) (b.1) (c.1)

(a.2) (b.2) (c.2)

(a.3) (b.3) (c.3)

Figure 7. Detection error rates of (a) multi-part fusion, (b) multi-sample fusion and (c) multi-part and multi-sample fusion for noisy/poor quality samples (glasses, eyelid obstruction and gaze deviation)

Results obtained when integrating the multi-part and multi-sample schemes are shown in Figure 7 c.1-c.3. It is clear that in the three cases under study, the trade-off between detection errors can be controlled by changing the number of parts and samples, decreasing the FRR and improving robustness against the attacks. Apart from that, lower error rates can be achieved than in the case of the non-fused base algorithm using the whole iris. It is necessary to take into account, however, that since samples are chosen randomly from the test data set results change from execution to execution. To guarantee that results obtained are statistically significant, average error rates and standard deviations have been estimated by repeatedly selecting random sets of *'m'* samples. Such results are presented in Table 3 for the best selection of parts and samples. The selection criteria is mainly based in minimum total error rate achieved, however, if highly similar results are obtained for different numbers of parts and samples, the case with the minimum number of them is considered provided that the standard deviation does not increase considerably.

Table 3. Noise/poor quality: average error rates (%) with standard deviation for best cases

Degradation	Error rate (%)	Whole iris (reference)	Best selection of parts and samples (m,n): (parts, samples)	
Glasses	FAR	0.746 $^{\pm 0.18}$	1/2 iris (2,3)	0.512 $^{\pm 0.09}$
			1/3 iris (2,3)	0.716 $^{\pm 0.19}$
			1/4 iris (3,3)	0.307 $^{\pm 0.08}$
			1/6 iris (3,3)	0.244 $^{\pm 0.08}$
	FRR	2.130 $^{\pm 2.73}$	1/2 iris (2,3)	0 $^{\pm 0}$
			1/3 iris (2,3)	0.087 $^{\pm 0.61}$
			1/4 iris (3,3)	0.087 $^{\pm 0.61}$
			1/6 iris (3,3)	0 $^{\pm 0}$
Eyelid obstruction	FAR	4.016 $^{\pm 0.87}$	1/2 iris (2,3)	2.263 $^{\pm 0.51}$
			1/3 iris (2,3)	2.564 $^{\pm 0.40}$
			1/4 iris (3,3)	1.066 $^{\pm 0.29}$
			1/6 iris (3,3)	1.587 $^{\pm 0.34}$
	FRR	8.005 $^{\pm 5.98}$	1/2 iris (2,3)	4.461 $^{\pm 4.53}$
			1/3 iris (2,3)	0.385 $^{\pm 1.68}$
			1/4 iris (3,3)	0.154 $^{\pm 1.08}$
			1/6 iris (3,3)	0.846 $^{\pm 2.42}$
Gaze deviation	FAR	7.8246 $^{\pm 1.12}$	1/2 iris (2,2)	5.215 $^{\pm 0.87}$
			1/3 iris (3,3)	7.864 $^{\pm 1.12}$
			1/4 iris (3,3)	8.487 $^{\pm 1.03}$
			1/6 iris (4,3)	6.033 $^{\pm 0.80}$
	FRR	9.400 $^{\pm 6.77}$	1/2 iris (2,2)	6.467 $^{\pm 5.39}$
			1/3 iris (3,3)	6.133 $^{\pm 5.58}$
			1/4 iris (3,3)	6.000 $^{\pm 4.59}$
			1/6 iris (4,3)	13.33 $^{\pm 5.27}$

When analysing in detail the results shown in Table 3, it can be noticed that regardless of the parts size (rings width), no more than 3 parts are usually required to achieve the best result possible. This fact is quite promising, since avoiding the use of the whole iris has some advantages like efficiency improvement or protection of the trait in a cancellable iris biometrics scenario. In the case of glasses and eyelid obstruction, better results that are statistically significant are obtained for all possible part sizes. In the case of using glasses, the main reason for this to happen is that different samples are usually affected in a different way by glasses (e.g. a slight head tilt can noticeably change the position of reflections), so using several samples helps minimizing the effect of the degradation. Using only certain iris parts can also help, especially if those parts affected by reflections are not used. Something similar occurs with iris obstruction. In the case of gaze deviation, worse results are obtained and lower error rates than the reference cannot always be achieved – see e.g. error rates when using fourths and sixths of the iris. When iris is off-axis all parts are affected (the most affected region depends on the deviation direction) and it is easy that all samples also are, so good results cannot be guaranteed just by using this architecture.

B. Artificially provoked iris alterations

When dealing with artificially provoked iris alterations, the behaviour of the individual multi-instance and multi-sample schemes is the same as in the case of noisy or poor quality samples, as it can be observed in Figure 8 a.1-a.2 and b.1-b.2. Results obtained when integrating the multi-part and multi-sample schemes are shown in Figure 8 c.1-c.2 and complemented by values in Table 4.

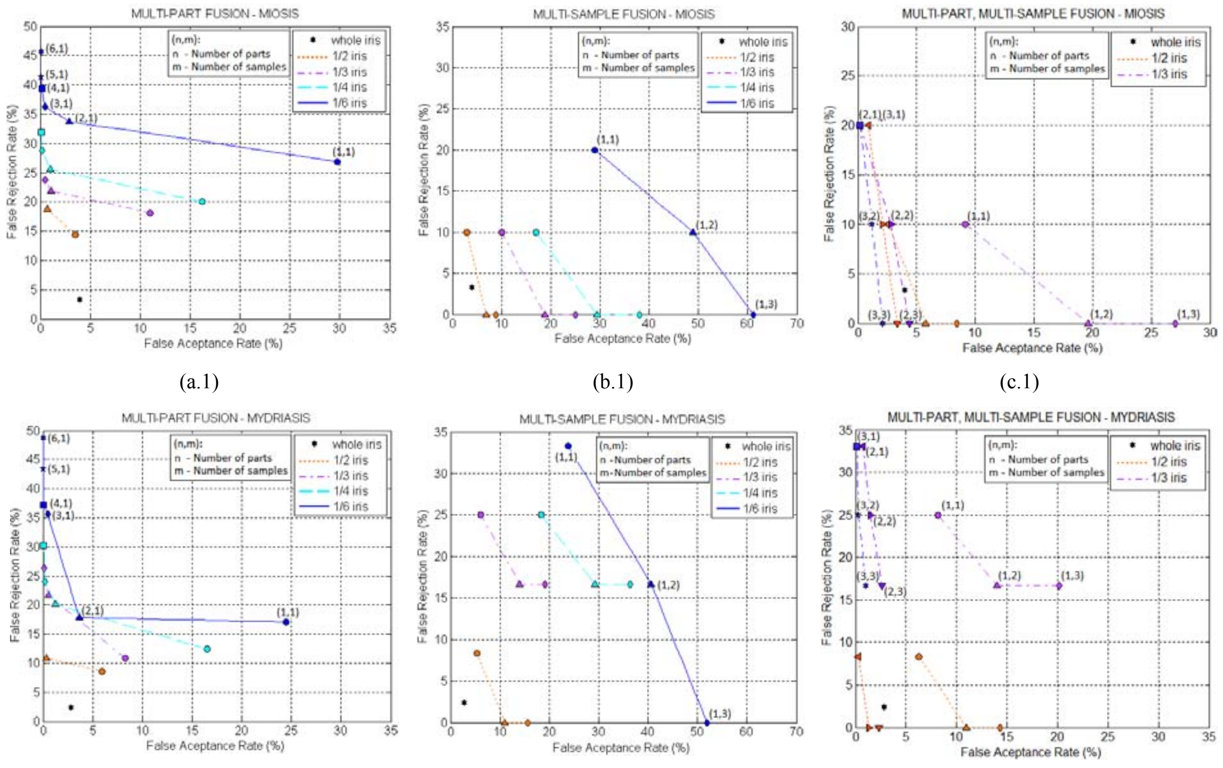

(a.1) (b.1) (c.1)

(a.2) (b.2) (c.2)

Figure 8. Detection error rates of (a) multi-part fusion, (b) multi-sample fusion and (c) multi-part and multi-sample fusion for artificially provoked iris alterations (miosis and mydriasis)

Table 4. Artificially provoked alterations: average error rates (%) with standard deviation for best cases

Degradation	Error rate (%)	Whole iris (reference)	Best selection of parts and samples (m,n): (parts, samples)	
Miosis	FAR	$4.235^{\pm 1.23}$	1/2 iris (2,3)	$3.446^{\pm 0.73}$
			1/3 iris (2,3)	$4.219^{\pm 1.01}$
			1/4 iris (3,3)	$2.153^{\pm 0.53}$
			1/6 iris (4,3)	$1.817^{\pm 0.39}$
	FRR	$5.800^{\pm 6.38}$	1/2 iris (2,3)	$1.200^{\pm 3.27}$
			1/3 iris (2,3)	$1.700^{\pm 3.77}$
			1/4 iris (3,3)	$1.900^{\pm 3.94}$
			1/6 iris (4,3)	$3.500^{\pm 5.00}$
Mydriasis	FAR	$2.073^{\pm 0.48}$	1/2 iris (2,3)	$2.557^{\pm 0.42}$
			1/3 iris (3,3)	$1.044^{\pm 0.32}$
			1/4 iris (3,3)	$1.246^{\pm 0.32}$
			1/6 iris (4,3)	$1.181^{\pm 0.27}$
	FRR	$2.750^{\pm 4.11}$	1/2 iris (2,3)	$0.417^{\pm 1.82}$
			1/3 iris (3,3)	$11.83^{\pm 3.67}$
			1/4 iris (3,3)	$20.83^{\pm 4.35}$
			1/6 iris (4,3)	$21.66^{\pm 6.37}$

Comparing the multi-part, multi-sample results with the reference, it is clear that robustness against this type of obfuscation attacks can also be achieved, with lower error rates that are statistically significant, especially in the case of miosis. In the case of mydriasis, this is only possible when using at least half of the iris as each of the parts. Preliminary results show that for mydriasis, better results can be achieved without using the whole iris when using the second and third rings of the iris (out of 3) instead of the first and second rings. The reason why this happens is that the non-elastic deformations of the iris when the pupil excessively dilates degrade most severely the ring nearest to the pupil. Thus, ring order can also be considered to further improve the results.

C. Occlusion of the iris

The last case to be analysed is the case of partial and total iris occlusion when using lenses. When occlusion is partial, the behaviour of the individual multi-instance and multi-sample schemes is the same as the one observed up to now. For total occlusion, results also seem to follow the expected trend; however, some incongruities can be observed like the fact that smaller parts of the iris provoke noticeably less false rejection errors than bigger parts (see results for sixths and halves of the iris in Figure 9 a.2). This incongruent behaviour can also be perfectly observed in the case of the multi-part and multi-sample integrated scheme in Figure 9 c.2. Results obtained when using a higher number of parts are worse than those obtained when using less. Since cosmetic lenses are opaque, the corresponding iris texture becomes unavailable, making recognition unfeasible. In this regard, the poor results obtained were expected.

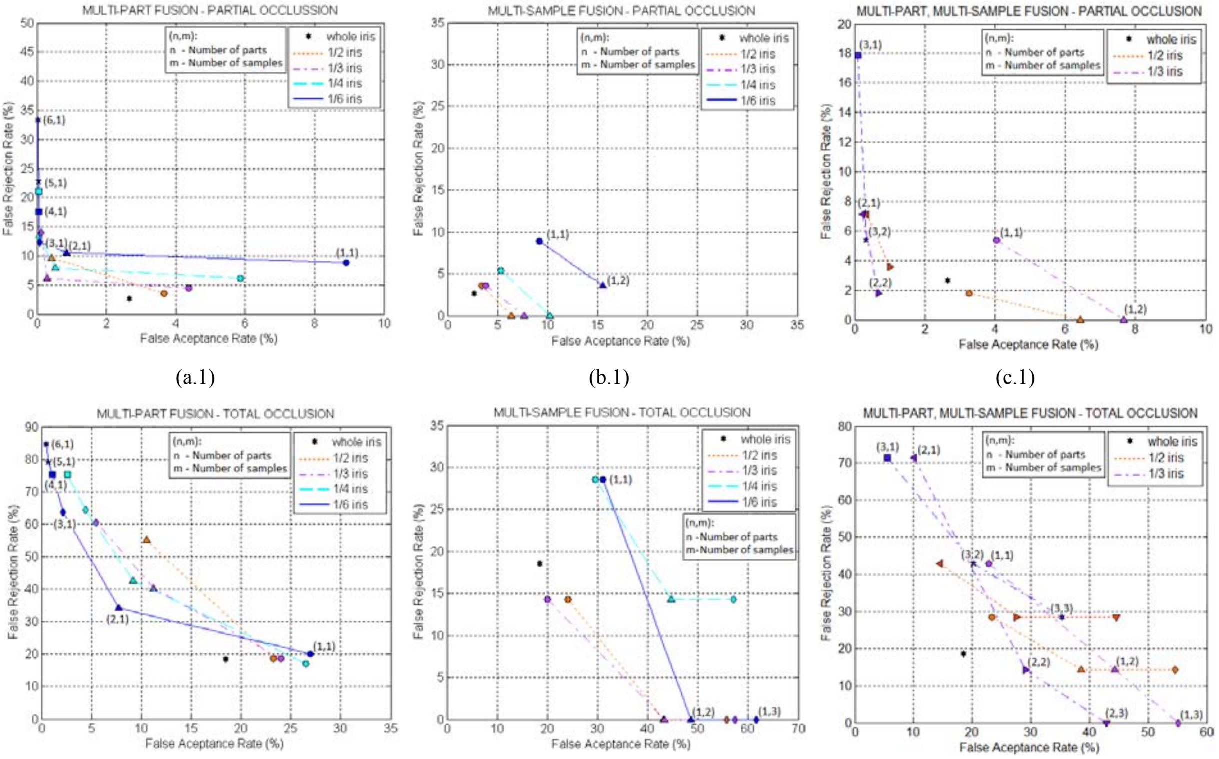

(a.1) (b.1) (c.1)

(a.2) (b.2) (c.2)

Figure 9. Detection error rates of (a) multi-part fusion, (b) multi-sample fusion and (c) multi-part and multi-sample fusion for occlusion (partial and total)

Table 5 shows the error rates obtained for partial occlusion for the best selection of parts and samples. As in previous cases, it can be noticed that regardless of the parts size, no more than 3 parts are required to achieve good results. Since only two samples are available from the database and both of them are used, standard deviation of the data is 0. Using only two samples does not allow enough control over the trade-off between FAR and FRR, and better results seem likely to be achieved by considering at least one more sample (m=3). In spite of this fact, better results than in the reference case are achieved for halves and fourths of the iris.

Although part of the iris texture is unavailable in this case, as long as there is enough usable information to be processed, obtaining such results is possible.

Table 5. Occlusion: average error rates (%) with standard deviation for best cases

Degradation	Error rate (%)	Whole iris (reference)	Best selection of parts and samples (m,n): (parts, samples)	
Partial occlusion	FAR	$2.578^{\pm 0.18}$	1/2 iris (2,2)	$1.002^{\pm 0}$
			1/3 iris (2,2)	$0.689^{\pm 0}$
			1/4 iris (2,2)	$1.378^{\pm 0}$
			1/6 iris (3,2)	$2.224^{\pm 0}$
	FRR	$2.696^{\pm 1.46}$	1/2 iris (2,2)	$3.571^{\pm 0}$
			1/3 iris (2,2)	$1.786^{\pm 0}$
			1/4 iris (2,2)	$1.786^{\pm 0}$
			1/6 iris (3,2)	$3.571^{\pm 0}$
Total occlusion	FAR	$17.604^{\pm 3.27}$	N/A	
	FRR	$25.429^{\pm 14.01}$	N/A	

5. Conclusions

Iris recognition systems are vulnerable to sample presentation attacks of the obfuscation type that increase false rejection rates. In this paper a multi-part, multi-sample sequential decision fusion architecture is applied to an iris recognition system to reduce the effect of such attacks. The system is tested with real data for degradations such as miosis, mydriasis, glasses, eyelid obstruction, gaze deviation and occlusion with lenses. The proposed architecture is demonstrated to provide robustness under obfuscation attacks with lower error rates and better control over the trade-off between FAR and FRR.

References

[1] P. J. Grother, G. W. Quinn, J. R. Matey, M. L. Ngan, W. J. Salamon, G. P. Fiumara and C. I. Watson, “IREX III - Performance of Iris Identification Algorithms”, *NIST Interagency/Internal Report (NISTIR)*, vol. 7836, 2012.

[2] V. P. Nallagatla and V. Chandran, “Sequential decision fusion for controlled detection errors”, *International Conference on Information Fusion*, 2010.

[3] V. P. Nallagatla and V. Chandran, “Sequential fusion using correlated decisions for controlled verification errors”, *International Conference on Computer Analysis of Images and Patterns*, vol. 2, pp. 49-56, 2011.

[4] T. Joshi, S. Dey and D. Samanta, “Multimodal biometrics: state of the art in fusion techniques”, *International Journal of Biometrics*, vol. 1, pp. 393-417, 2009.

[5] A. A. Ross, K. Nandakumar and A. K. Jain, *Handbook of multibiometrics*, Springer, 2006.

[6] A. K. Jain, A. A. Ross and K. Nandakumar, *Introduction to biometrics*, Springer, 2011.

[7] N. K. Ratha, J. H. Connell and R. M. Bolle, “Enhancing security and privacy in biometrics-based authentication systems”, *IBM Systems Journal*, vol. 40, no.3, pp.614-635, 2001.

[8] A. K. Jain, S. Pankanti, S. Prabhakar, L. Hong and A. A. Ross, “Biometrics: A grand challenge”, *International Conference on Pattern Recognition*, 2004.

[9] J.L. Wayman, “Technical testing and evaluation of biometric devices”, *Biometrics - Personal Identification in Networked Society*, Kluwer Academic Publisher, 1999.

[10] B. Cukic and N. Bartlow, “The vulnerabilities of biometric systems - An integrated look and old and new ideas”, *Technical report*, 2005.

[11] Common Criteria, *Common Methodology for Information Technology Security Evaluation (CEM v3.1)*, 2009.

[12] Common Criteria, *Biometric Evaluation Methodology Supplement (BEM)*, 2002.

[13] Z. Geradts and P. Sommer, “Forensic implications of identity management systems”, *FIDIS NoE WP6 Deliverable*, 2006.

[14] S. A. C. Schuckers, “Spoofing and anti-Spoofing measures”, *Information Security Technical Report*, vol. 7, no. 4, pp. 56-62, 2002.

[15] I. Tomeo-Reyes, J. Liu-Jimenez, I. Rubio-Polo, J. Redondo- Justo and R. Sanchez-Reillo, “Input images in iris recognition systems: A case study”, *IEEE International Systems Conference*, pp. 501-505, 2011.

[16] H. Proença and L. A. Alexandre (Eds.), “Noisy Iris Challenge Evaluation II - Recognition of Visible Wavelength Iris Images Captured At-a-distance and On-the-move”, *Pattern Recognition Letters*, vol. 33, no.8, pp. 963-1026, 2012.

[17] K. Hollingsworth, K. Bowyer and P. Flynn, “Pupil dilation degrades iris biometric performance”, *Computer Vision and Image Understanding*, vol. 113, no.1, pp.150-157, 2009.

[18] J. G. Daugman, “High confidence visual recognition of persons by a test of statistical independence”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15, no.11, pp. 1148-1161, 1993.

[19] H. Patel, C. K. Modi, M. C. Paunwala and S. Patnaik, “Human identification by partial iris segmentation using pupil circle growing based on binary integrated edge intensity curve”, *International Conference on Communication Systems and Network Technologies,* pp.333-338, 2011.

[20] J. K. Pillai, V. M. Patel, R. Chellappa and N. K. Ratha, “Sectored random projections for cancelable iris biometrics”, *IEEE International Conference on Acoustics Speech and Signal Processing*, pp. 1838-1841, 2010.

[21] M. R. Islam, Y. C. Wang and A. Khatun, “Partial iris image recognition using wavelet based texture features”, *International Conference on Intelligent and Advanced Systems*, pp.1-6, 2010.

[22] C. C. Teo and H. T. Ewe, “An efficient one-dimensional fractal analysis for iris recognition”, *International Conference on Computer Graphics, Visualization and Computer Vision*, pp.157-160, 2005.

[23] J. Canny, “A Computational approach to edge detection”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no.6, pp. 679-698, 1986.

[24] R. Deriche, J. P. Cocquerez and G. Almouzni, “An efficient method to build early image description”, *International Conference on Pattern Recognition*, 1988.