Decision Fusion from Parts and Samples for Robust Iris Recognition

Inmaculada Tomeo-Reyes
Queensland University of Technology
2 George St., Brisbane, QLD, 4000, Australia
inma.tomeoreyes@qut.edu.au

Vinod Chandran
Queensland University of Technology
2 George St., Brisbane, QLD, 4000, Australia
v.chandran@qut.edu.au

Abstract

Fusion techniques can be used in biometrics to achieve higher accuracy. 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. In a fused system, varying a single threshold does not allow this to be achieved, but systematic adjustment of a set of parameters does. In this paper, fused decisions from a multi-part, multi-sample sequential architecture are investigated for that purpose in an iris recognition system. A specific implementation of the multi-part architecture is proposed and the effect of the number of parts and samples in the resultant detection error rate is analysed. The effectiveness of the proposed architecture is then evaluated under two specific cases of obfuscation attack: miosis and mydriasis. Results show that robustness to such obfuscation attacks is achieved, since lower error rates than in the case of the non-fused base system are obtained.

1. Introduction

Using fusion techniques to achieve higher accuracy in the biometric field is now common practice. Within the algorithm framework, considerable research has been done into fusion of information from multiple sources at different levels: data, feature, score and decision levels [1]. Matcher decisions are the smallest and most unambiguous piece of information for fusion. For this reason, decision level fusion allows single modality based systems to combine even if these systems have been developed independently, possibly by different companies, and internal workings are not made public. This is a very common situation 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 information sources. There are several approaches that can be considered at the decision level [1, 2]. Existing approaches include simple rules as the AND/OR rules, majority voting and weighted majority voting. The performance of these approaches is compared in [3, 4] for different multibiometric systems. Other approaches include Bayesian fusion rules [5] or the Behaviour Knowledge Spaced (BSK) method [6]. The Dempster-Shafer theory of evidence was used by Kuncheva in [7], and compared with more than ten different classifier fusion techniques.

Recent investigations by Nallagatla and Chandran [8, 9] using text-dependent speech have shown that in a sequential decision fusion architecture with multiple instances and multiple samples, it is possible to control the trade-off between false accept and false reject errors using the number of instances and samples considered in the architecture. Such control is not guaranteed by other methods of classifier combination such as adaptive boosting [10], for example, where the objective is to obtain optimal performance (least total error). In this case, thresholds at each classifier stage are set accordingly and once the chain of classifiers is selected there are no parameters that control the trade-off between errors. Controlled trade-off is especially desirable when biometric systems are in operation and the threat level changes. In general, under normal operating conditions false acceptances are kept very low and false rejections may be reasonably higher. Any false rejections are usually processed with less reliable and less secure means as fallback mechanisms. However, a person on a watch list can take advantage of this and use a sample presentation attack such as alteration of his iris or face to prevent being matched to the watch list. In these situations it is desirable to lower false rejections. Additional biometric information may then be called into operation when the threat level (or probability of such attacks) is considered high. For such scenarios, it is necessary to understand how biometric fusion can be used to achieve desired objectives.

The adaptation and application of a multi-instance, multi-sample sequential decision fusion architecture to iris recognition in the context of preventing sample presentation attacks is a research problem that has not been addressed yet. This paper is aimed at investigating the viability of such an approach in detail. How sample presentation attacks, and more specifically, obfuscation attacks, affect iris recognition systems is explained in section 2. Section 3 provides details of the fusion architecture. Section 4 describes the methodology used to evaluate the fused system performance and the results. Finally, a brief conclusion with suggestions for potential future work is included in section 5.

2. Sample presentation attacks

Any biometric system, regardless of the trait, is basically composed of four different subsystems: data acquisition, pre-processing, feature extraction and comparison. Each subsystem may have different points of attack, and for each point of attack, there may be one or more potential exploits. Some of the early work by Ratha et al. [11] identified eight possible points of attack. Further work by Jain et al. [12] and Wayman [13] sought to refine this approach. Bartlow and Cukic [14] and Common Criteria [15] extended this research by adding administrative components.

Among all potential attacks, this paper focuses on attacks at the user interface level or sample presentation attacks, which are defined as any attempt to break into a system by presenting a biometric sample. These attacks can be categorized into three main groups [16]: *impersonation* (an impostor attempts to intrude the system by posing himself as another authorized user), *obfuscation or disguise* (attacker deliberately changes his/her biometric characteristic to avoid being recognized) and *spoofing* (a counterfeit biometric trait that is not obtained from a live claimant is used to achieve positive recognition). 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 in these attacks in this paper. Threats associated with *obfuscation attacks* in iris recognition can be grouped into three categories:

A. Intentional presentation of noisy or poor-quality samples

A noisy, poor-quality or null iris sample that may not match the attacker template can be intentionally presented to the system to avoid recognition. Blinking, deviating the gaze and using glasses are typical ways to achieve this.

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 [17]; however, this research did not extend to extreme cases of pupil dilation or constriction 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. Where they are opaque, they totally occlude the iris, making the corresponding iris texture unavailable.

It is a fact that the above-mentioned attacks degrade iris recognition performance [18], regardless of its high accuracy. In order to demonstrate this, an iris recognition algorithm inspired by Daugman's works [19] has been used here. The segmentation process combines the black hole search method [20] and a simplified version of Daugman's integro-differential operator. Feature extraction is based on Gabor filters and Hamming distance is used for classification (see subsection 4.1). 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 1). Results clearly illustrate the false rejections increment typical from obfuscation attacks.

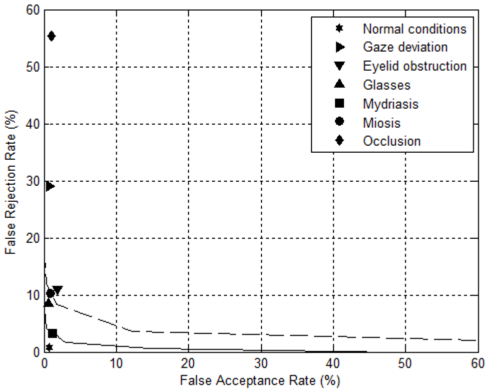

Figure 1: Detection error rates for different obfuscation attacks considering system threshold under normal operating conditions

A multi-part, multi-sample architecture for robust iris recognition is proposed in this research work. To demonstrate that such architecture can improve the performance of an iris recognition system under the threat of an obfuscation attack, artificially provoked iris alterations (mydriasis and miosis) are considered. The objective in this work is to deal with and minimize the impact of samples affected by severe mydriasis and miosis, instead of just detecting and discarding the sample, as commonly done in currently deployed systems. For the two attacks chosen, DET curves have been added to Figure 1 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.

3. Multi-biometric decision fusion techniques

3.1. Multi-part fusion scheme

For iris recognition, using both irises as different instances in a multi-classifier or multi-instance fusion scheme is the most common approach. However, this option does not seem to be useful in the case of obfuscation attacks, since if the impostor deliberately changes his/her iris to avoid being recognized, such modification can most probably be done for both irises. Instead, using different parts of the same iris image as the different instances seems to be a reasonable approach for various reasons. First of all, 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. Secondly, both under normal and threatening conditions, noise sources as eyelids or bright spots can be avoided. In general, iris recognition using parts has been considered in the literature for different purposes, such as dealing with noise [21], cancellable iris biometrics [22] or efficiency improvement [23]. Here, the part based approach is mainly used to make the base architecture robust to obfuscation attacks, although some of the above-mentioned advantages are also applicable.

There are many schemes for dividing an iris into parts. Concentric rings of equal width have been chosen in this paper. Iris templates are built for each ring and classifiers for the different rings form the stages of the sequential decision fusion architecture. Only if the decisions from all stages are *accept*, the claim is accepted. This is equivalent to applying an AND decision fusion rule. This fusion method effectively reduces the false acceptances; however, it can increase the number of false rejections. In the case of a multi-part scheme with '*n*' stages, the False Acceptance Rate (FAR) and False Rejection Rate (FRR) of the fused system in terms of false acceptance rate 'α' and false rejection rate 'ρ' for each statistically independent classifier can be calculated as follows [8]:

$$\alpha_{fused} = \alpha^n \tag{1}$$

$$\rho_{fused} = \rho + (1-\rho)\rho + (1-\rho)^2\rho + \cdots + (1-\rho)^{n-1}\rho \tag{2}$$

Equation (2) can be reduced to $\rho_{\text{fused}} \approx n\rho$ when $\rho << 1$, but such simplification is not appropriate when dealing with obfuscation attacks, since false rejections are significantly increased with respect to normal operating conditions.

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. At any given stage (part classifier), there are multiple samples that can be presented if the decision is *reject*. The number of samples is limited to a maximum allowable. If this number 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. This fact becomes clear when considering equations (3) and (4) [8], which allow the calculation of the FAR and FRR of a multi-sample fusion scheme with '*m*' maximum repeated attempts in terms of the probabilities of false acceptance and false rejection for each independent attempt (α and ρ).

$$\alpha_{fused} = m\alpha \tag{3}$$

$$\rho_{fused} = \rho^m \tag{4}$$

3.3. Multi-part and multi-sample schemes integration

The resulting architecture after combining the multi-part and multi-sample schemes can be seen in Figure 2. It includes '*n*' classifiers arranged sequentially and allows up to '*m*' samples presentation (dashed line) in case any of the user's claims is not accepted by the system (d=0).

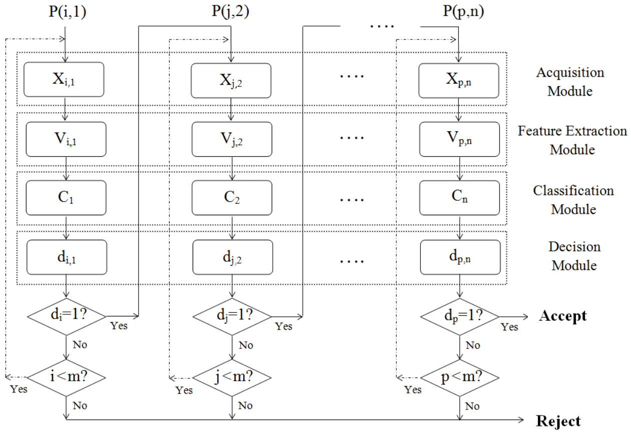

Figure 2: Multi-part, multi-sample decision fusion scheme

From equations (1)-(4), FAR and FRR of the fused system in terms of false acceptance and false rejection for each independent attempt (α and ρ) are calculated according to (5)-(6). 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.

$$\alpha_n^m = (m\alpha)^n \tag{5}$$

$$\rho_n^m = \rho^m + (1-\rho^m)\rho^m + \cdots + (1-\rho^m)^{n-1}\rho^m \tag{6}$$

4. Experimental validation

4.1. Iris recognition base algorithm

The implemented iris recognition base algorithm is based on Daugman’s approach [19], although it includes some modifications. In the pre-processing stage, the black hole search method [20] is first used to locate the pupil (course search). A simplified version of Daugman’s integro-differential operator is then used (fine search) to define 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 [24] and Deriche et al. algorithm [25] is used. Regarding 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 after barrel-shifting to prevent iris rotation. As a result, a 3072 bits iris code is obtained. 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%, close to other available baselines that represent the current state of the art [26].

4.2. Database

In the case of artificially provoked iris alterations, a miotic/mydriatic agent in the form of eyedrops was instilled to participants to create the dataset, constituted by a total of 654 images (see Table 1) acquired at a resolution of 640x468. Although the initial dataset size was larger, around 30% of the images were discarded because of segmentation errors. The increase in segmentation errors is an important effect of miosis and mydriasis 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. 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.

Table 1. Iris dataset: artificially provoked iris alterations

	Users	Images	Train set	Test set
Normal conditions	59	177	59	118
Mydriasis	14	229	42	187
Miosis	10	248	30	218

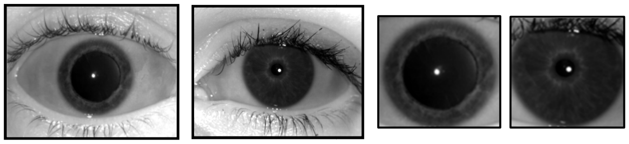
(a.1) (b.1) (a.2) (b.2)

Figure 3: Sample images with details. (a) Mydriasis (b) Miosis

4.3. Train and test strategy

In order to design the architecture and validate the results, a train subset is first used to fix the threshold of each of the single classifiers, this being the EER-based threshold. Results are later obtained using unseen images from the test subset (see Table 1). The train and test subsets are disjoint. The EER from the train subset when considering the whole iris is used as a reference. It may be noted that the same analysis can be performed using any classifier threshold different from the EER-based one (e.g. a threshold that allows to keep false acceptances very low, as usual in commercial systems). 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 clarify which values belong to the same experiment. To calculate the error rates, all tests match samples against templates captured under normal conditions (59 templates in total). Each user with available samples of a given degradation (e.g. 14 in the case of mydriasis) is selected for the genuine tests. Each degraded sample of such a user is matched against all other users (58 in total) for impostor tests.

4.4. Results

A. Multi-part fusion results.

Results for multi-part fusion are obtained by progressively increasing the number ‘n’ of rings or parts. Four different values of ‘n’ are considered (n = 2, 3, 4 and 6), being the width of each ring equal to the width of the whole normalized iris divided by ‘n’ (only the usable iris area is considered for normalization). 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. According to equations (1)-(2), when considering multi-part fusion, the FAR decreases multiplicatively with the number ‘n’ of classifiers or parts, while the FRR increases as the addition of the term

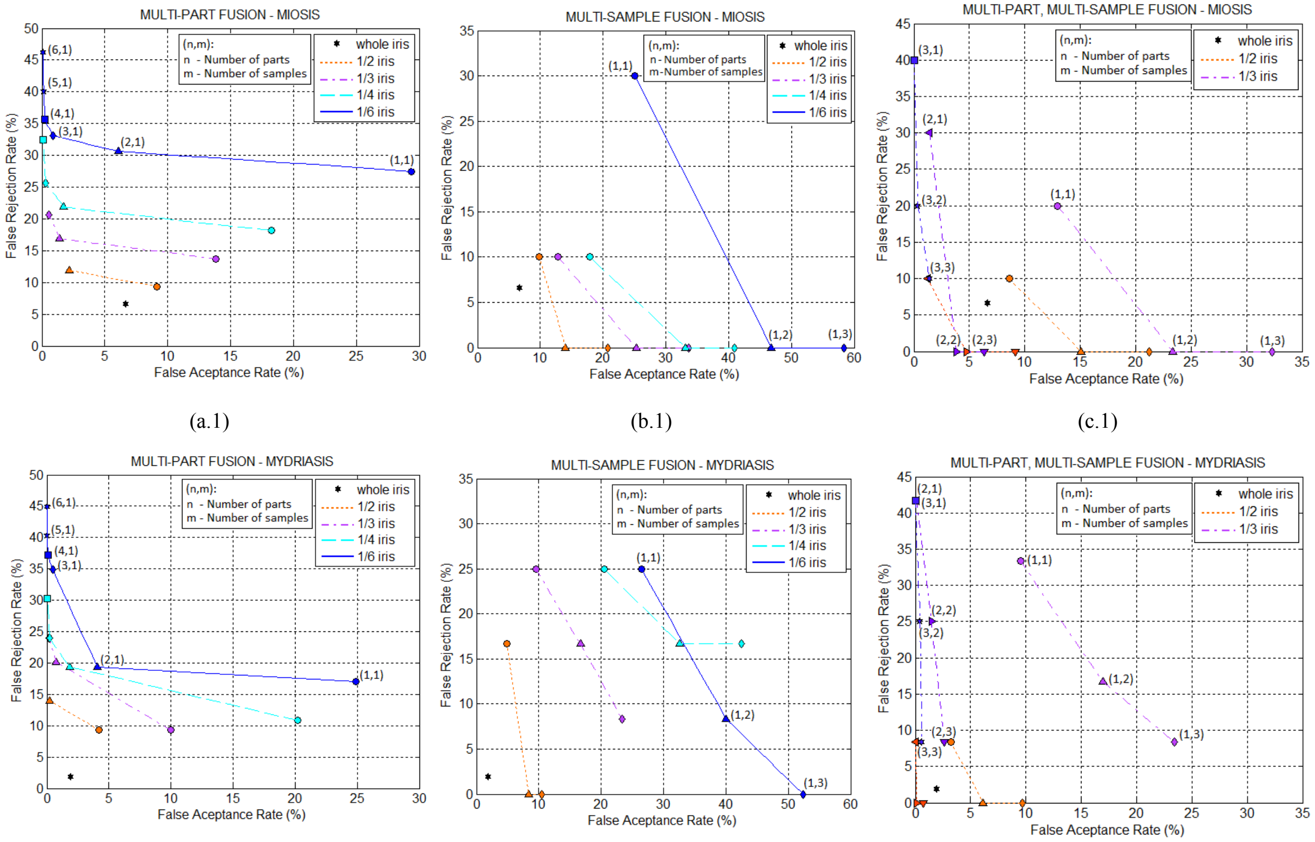

(a.1) (b.1) (c.1)

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

Figure 4: Detection error rates of (a) multi-part fusion, (b) multi-sample fusion and (c) multi-part and multi-sample fusion for miosis and mydriasis. The EER for the whole iris in each case (miosis/mydriasis) is used as a reference (see star symbol in each Figure).

$(1-\rho)^{n-1}\rho$ from 1 to 'n'. Since multiplicative changes are faster than additive ones, the reduction in the FAR is faster than the increase in the FRR, especially if the number of parts used is small. This behaviour can be clearly observed in Figure 4 a.1-a.2. Results also show that by adding enough stages to the architecture FAR can be made low at the cost of a higher FRR. This trend is not useful with obfuscation attacks.

B. Multi-sample fusion results.

In the case of multi-sample, a maximum of three sample presentations are allowed to verify the user (m = 3). Samples are chosen randomly from the test data set. 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). Results for an increasing number 'm' of samples (m = 1, 2 and 3) are shown in Figure 4 b.1-b.2. As it can be observed, the more samples considered the lower value of the FRR, but since the FAR increases, it is clear that for the total error to decrease, a combination of multiple parts and samples is needed. Consistent with equation (4), the FRR is decreasing multiplicatively with the number 'm' of attempts. However, the FAR does not increase exactly in an additive way, as stated in equation (3). It may be noted that equations (3)-(4) assume statistical independence, which is usually an unrealistic assumption. Methods to select a specific subset of classifiers to achieve optimal performance considering statistical dependence will be explored as future work.

C. Multi-part, multi-sample fusion results.

Results obtained when integrating the multi-part and multi-sample schemes are shown in Figure 4 c.1-c.2. For clarity purposes, just the cases in which halves and thirds of the iris (n = 2, 3) are used as parts are considered. Table 2 shows the results obtained with the best selection of parts and samples among those represented in Figure 4 c.1-c.2. Comparing the multi-part, multi-sample results with the reference (best error rate achieved from non-fused base algorithm with the whole iris), it can be observed that robustness against the obfuscation attacks is achieved, with lower error rates that are statistically significant. Even more promising is the fact that, in some cases, it is not necessary to use the whole iris for that. This can be observed for example in the case of miosis, when using 2 out of 3 thirds of the iris and 3 samples.

Table 2. Average error rates with standard deviation for best cases

		Whole iris (reference)	Best selection of parts and samples	
Miosis	FAR	$5.189^{\pm 1.20}$	1/3 iris (2,2)	$3.643^{\pm 0.96}$
			1/3 iris (2,3)	$6.398^{\pm 1.26}$
	FRR	$6.800^{\pm 7.12}$	1/3 iris (2,2)	$4.706^{\pm 6.11}$
			1/3 iris (2,3)	$1.176^{\pm 3.25}$
Mydriasis	FAR	$1.266^{\pm 0.42}$	1/2 iris (2,3)	$1.085^{\pm 0.35}$
	FRR	$2.666^{\pm 4.2}$	1/2 iris (2,3)	$0.833^{\pm 2.52}$

5. Conclusion and future work

A multi-part, multi-sample sequential decision fusion architecture is applied to an iris recognition system and demonstrated to provide robustness under obfuscation attacks (mydriasis and miosis) by (a) controlled error trade-off and (b) lower error rates. 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 improve the results. Future work will investigate the effects of ordered parts and statistical dependence between classifier decisions.

Acknowledgments

The authors acknowledge Prof. David Atchison for his assistance with data collection. This work has been partly funded by National Security Science and Technology (NSST).

References

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

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

[3] K. A. Toh and W. Y. Yau. Combination of hyperbolic functions for multimodal biometrics data fusion, IEEE Trans. on Systems, Man and Cybernetics, 34 (2):1196-1209, 2004.

[4] M. M. Monwar and M. Gavrilova. A robust authentication system using multiple biometrics, Computer and Information Science, Springer, 131:189-201, 2008.

[5] K. Veeramachaneni, L. A. Oscadciw and R. K. Varshney. An adaptive multimodal biometric management algorithm, IEEE Transactions on Systems, Man and Cybernetics, Part A, 35(3):344-356, 2005.

[6] F. Roli, G. Fumera and J. Kittler. Fixed and trained combiners for fusion of imbalanced pattern classifiers, Int'l Conference on Information Fusion, 1:278-284, 2002.

[7] L. L. Kuncheva, J. C. Bezdek and R. P. W. Duin. Decision templates for multiple classifier fusion: An experimental comparison, Pattern Recognition, 34(2):299-314, 2001.

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

[9] V. P. Nallagatla and V. Chandran. Sequential fusion using correlated decisions for controlled verification errors, International Conference on Computer Analysis of Images and Patterns, 2:49-56, 2011.

[10] Y. Freund and R. E. Schapire. A desicion-theoretic generalization of on-line learning and an application to boosting, Computational learning theory, Springer, 1995.

[11] N. K. Ratha, J. H. Connell and R. M. Bolle. Enhancing security and privacy in biometrics-based authentication systems, IBM Systems Journal, Vol. 40(3), 2001.

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

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

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

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

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

[17] K. Hollingsworth, K. Bowyer and P. Flynn. Pupil dilation degrades iris biometric performance, Computer Vision and Image Understanding, 113(1):150–157, 2009.

[18] 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, 501-505, 2011.

[19] J. G. Daugman. High confidence visual recognition of persons by a test of statistical independence, IEEE Transactions on Pattern Analysis and Machine Intelligence, 15(11):1148-1161, 1993.

[20] C.C. Teo and H.T. Ewe. An efficient one-dimensional fractal analysis for iris recognition. Int'l Conference on Computer Graphics,Visualization and Computer Vision,157-160,2005.

[21] 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, 333-338, 2011.

[22] 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, 1838-1841, 2010.

[23] M. R. Islam, Y. C. Wang and A. Khatun. Partial iris image recognition using wavelet based texture features, Int'l Conference on Intelligent and Advanced Systems, 1-6, 2010.

[24] J. Canny. A Computational approach to edge detection, IEEE Transactions on Pattern Analysis and Machine Intelligence, 8(6):679–698, 1986.

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

[26] 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) 7836, 2012.