Secure and Robust Iris Recognition using Random Projections and Sparse Representations

Jaishanker K. Pillai, *Student Member, IEEE,* Vishal M. Patel, *Member, IEEE,* Rama Chellappa, *Fellow, IEEE,* and Nalini K. Ratha, *Fellow, IEEE.*

Abstract—Non-contact biometrics such as face and iris have additional benefits over contact based biometrics such as fingerprint and hand geometry. However, three important challenges need to be addressed in a non-contact biometrics-based authentication system: ability to handle unconstrained acquisition, robust and accurate matching and privacy enhancement without compromising security. In this paper, we propose a unified framework based on random projections and sparse representations, that can simultaneously address all the three issues mentioned above in relation to iris biometrics. Our proposed quality measure can handle segmentation errors and a wide variety of possible artifacts during iris acquisition. We demonstrate how the proposed approach can be easily extended to handle alignment variations and recognition from iris videos, resulting in a robust and accurate system. The proposed approach includes enhancements to privacy and security by providing ways to create cancelable iris templates. Results on public datasets show significant benefits of the proposed approach.

Index Terms—Iris Recognition, Cancelability, Secure Biometrics, Random Projections, Sparse Representations.

I. INTRODUCTION

Iris recognition is one of the most promising approaches for biometric authentication [1]. Existing algorithms based on extracting and matching features from iris have reported very high recognition rates on clean datasets [2]. However, since these methods rely on the fine texture features extracted from the iris, their performances degrade significantly when the image quality is poor [1], [3]. This seriously limits the application of the iris recognition system in practical scenarios, where the acquired image could be of low quality due to motion, partial co-operation or the distance of the user from the scanner.

When the acquisition conditions are not constrained, many of the acquired iris images suffer from defocus blur, motion blur, occlusion due to the eyelids, specular reflections and segmentation errors. Fig. 1 shows some of these distortions on images from the ICE2005 dataset [3]. Hence, it is essential to first select the "recognizable" iris images before employing the recognition algorithm. Recently, Wright *et al.* [4] introduced a sparse representation-based face recognition algorithm, which outperforms many state of the art algorithms when sufficient

number of training images are available. In this paper, we analyze the use of sparse representations for selection and recognition of iris images. We extend the original method [4] through a Bayesian fusion framework, where different sectors of the iris images are recognized separately using the above mentioned technique and the results of different sectors are combined based on their quality. This is significantly faster than the original method as it facilitates parallelization and reduces the size of the dictionary size, as will become apparent.
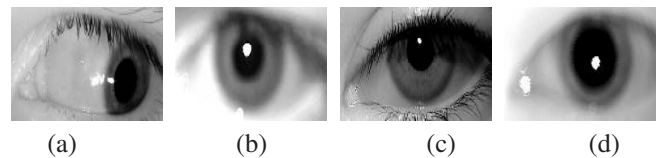

Fig. 1. Some poorly acquired iris images from the ICE dataset [3]. Note that image (a) has specular reflections on the iris and is difficult to be segmented correctly due to the tilt and non circular shape. Images (b) and (d) suffer from blurring, whereas image (c) is occluded by the shadow of the eyelids.

The performance of most existing iris recognition algorithms depends strongly on the effectiveness of the segmentation algorithm. Iris image segmentation normally involves identifying the ellipses corresponding to pupil and iris, and detecting the region inside these ellipses that is not occluded by the eyelids, eyelashes and specular reflections. Unfortunately, in unconstrained scenarios, correctly segmenting the iris images is extremely challenging [5]. The proposed selection algorithm removes input images with poorly segmented iris and pupil ellipses. Furthermore, since the introduced recognition scheme is robust to small levels of occlusions, accurate segmentation of eyelids, eyelashes and specular reflections are no longer critical for achieving good recognition performance.

Another important aspect in iris biometrics is security and privacy of the users. When the texture features of one's iris are stored in a template dictionary, a hacker could possibly break into the dictionary and steal these patterns. Unlike credit cards, which can be revoked and reissued, biometric patterns of an individual cannot be modified. So, directly using iris features for recognition is extremely vulnerable to attacks. To deal with this, the idea of cancelable iris biometrics has been introduced in [6], [7], [8], which can protect the original iris patterns as well as revoke and reissue new patterns when the old ones are lost or stolen. In this paper, we introduce two methods for incorporating security into the proposed iris recognition system, namely, random projections and random permutations.

Our methods can issue a different template for each application based on the original iris patterns of the person, generate a new template if the existing one is stolen while retaining the original recognition performance. The representation prevents extraction of significant information about the original iris patterns from cancelable templates.

A. Organization of the Paper

The paper is organized as follows. In section II, we discuss some of the existing algorithms for iris image selection, recognition and cancelability. The theory of sparse representation is summarized in section III. The Bayesian fusion framework for selecting and recognizing iris images is described in IV. We extend our method to video-based iris recognition in section V and discuss how to handle alignment in section VI. Two schemes for introducing cancelability into our framework are proposed in VII. Experiments and results are presented on simulated and real iris images in section VIII. We briefly outline future work and conclude in section IX.

II. PREVIOUS WORK

In this section, we briefly describe some of the existing methods for iris recognition, image quality estimation and cancelability.

A. Iris recognition

The first operational automatic iris recognition system was developed by Daugman [9] in 1993, in which Gabor features were extracted from scale normalized iris regions and quantized to form a 2K bit iris code. The Hamming distance between the iris code of the test and the training iris images was used for recognition. Wildes [10] used Laplacian of a Gaussian filter at multiple scales to produce a template and used the normalized correlation as the similarity measure. In recent years, researchers have analyzed aspects like utilizing real valued features for recognition, developing alternate ways of obtaining the binary codes and combining multiple features. See [1] for an excellent survey of recent efforts on iris recognition.

Several studies have shown that accurate quality estimation can improve the performance either by rejecting the poor quality images or by fusing the quality information during matching [1], [11], [12]. Rejection of poor quality images is useful in numerous practical settings, as explained below -:

1) A human assisted recognition system in which a small fraction of the samples where the system has lower confidence is sent to a human expert or a more accurate recognition system.
2) An active acquisition system in which a new sample from the user is acquired if the originally captured one is of poor quality.
3) High security applications requiring very low false positives, where it is acceptable to reject a small number of poor quality samples and deny access to individuals rather than obtain wrong results and provide access to unauthorized individuals.

We review some of the existing iris image quality estimation schemes in iris literature below. Daugman used the energy of the high frequency components as a measure of blur [9]. Proenca and Alexandre trained a neural network to identify common noise degradations in iris images [13]. Zhu *et al.* used the wavelet coefficients to evaluate the quality of iris images [14]. The Fourier spectra of local iris regions was used by Ma *et al.* to characterize blur and occlusion [15]. Rakshit and Monro used the quality and position of specular reflections for selecting good quality images [16]. With the exception of Daugman's method, these algorithms are specialized for image selection, which requires a separate method for recognizing iris images. Also, these algorithms utilize some property of the iris image to measure image quality and cannot handle the wide variety of common artifacts such as specular reflections and occlusion. In contrast to these methods, the image quality measure introduced in this paper can handle segmentation errors, occlusion, specular reflections, and blurred images. The proposed method also performs both selection and recognition in a single step.

B. Iris Recognition from Videos

Though research in iris recognition has been extremely active in the past decade, most of the existing results are based on recognition from still iris images [17]. Multiple iris images have been used in the past to improve performance. Du *et al.* [18] demonstrated higher rank one recognition rates by using three gallery images instead of one. Ma *et al.* [19] also enrolled three iris images and averaged the three Hamming distances to obtain the final score. Krischen *et al.* [20] used the minimum of the three Hamming distance as the final score. Schmid *et al.* [21] demonstrated that fusing the scores using log likelihood ratio gave superior performance when compared to average Hamming distance. Liu *et al.* [22], Roy and Bhattacharya [23] used multiple iris images for training classifiers.

The distortions common in iris image acquisition like occlusion due to eyelids, eye lashes, blur, and specular reflections will differ in various frames of the video. So by efficiently combining the different frames in the video, the performance could be improved. Temporal continuity in iris videos was used for improving the performance by Hollingsworth *et al.* [17]. The authors introduced a feature level fusion by averaging the corresponding iris pixels and a score level fusion algorithm combining all the pairwise matching scores. Though averaging reduces the noise and improves the performance, it required images to be well segmented and aligned, which may often not be possible in a practical iris recognition system. We will introduce a quality based matching score that gives higher weight to the evidence from good quality frames, yielding superior performance even when some video frames are poorly acquired.

C. Cancelable iris biometrics

The concept of cancelable biometrics was first introduced by Ratha *et al.* in [7], [8]. A cancelable biometric scheme intentionally distorts the original biometric pattern through a

revocable and non-invertible transformation. The objectives of a cancelable biometric system are as follows [6]:

- Different templates should be used in different applications to prevent cross matching.
- Template computation must be non-invertible to prevent unauthorized recovery of biometric data.
- Revocation and reissue should be possible in the event of compromise, and
- Recognition performance should not degrade when a cancelable biometric template is used.

In [24], Hash functions were used to minimize the compromise of the private biometric data of the users. Cryptographic techniques were applied in [25] to increase the security of iris systems. In [26], error correcting codes were used for cancelable iris biometrics. A fuzzy commitment method was introduced in [27]. Other schemes have also been introduced to improve the security of iris biometric. See [6], [24], [25], [26], [27], [28] and the references therein for more details.

The pioneering work in the field of cancelable iris biometric was done by Zuo *et al.* [29]. They introduced four non-invertible and revocable transformations for cancelability. While the first two methods utilized random circular shifting and addition, the other two methods added random noise patterns to the iris features to transform them. As noted by the authors, the first two methods gradually reduce the amount of information available for recognition. Since they are essentially linear transformations on the feature vectors, they are sensitive to outliers in the feature vector that arise due to eyelids, eye lashes and specular reflections. They also combine the good and bad quality regions in the iris image leading to lower performance. The proposed random projections based cancelability algorithm works on each sector of the iris separately, so outliers can only affect the corresponding sectors and not the entire iris vector. Hence, it is more robust to common outliers in iris data when compared to [29].

III. IRIS IMAGE SELECTION AND RECOGNITION

Following [4], in this section, we briefly describe the use of sparse representations for the selection of good quality iris images and their subsequent recognition.

A. Sparse Representations

Suppose that we are given L distinct classes and a set of n training iris images per class. We extract an N-dimensional vector of Gabor features from the iris region of each of these images. Let $\mathbf{D}_k = [\mathbf{x}_{k1}, \ldots, \mathbf{x}_{kj}, \ldots, \mathbf{x}_{kn}]$ be an $N \times n$ matrix of features from the k^{th} class, where $\mathbf{x}_{kj}$ denote the Gabor feature from the j^{th} training image of the k^{th} class. Define a new matrix or dictionary $\mathbf{D}$, as the concatenation of training samples from all the classes as

$$\mathbf{D} = [\mathbf{D}_1, ..., \mathbf{D}_L] \in \mathbb{R}^{N \times (n.L)}$$
$$= [\mathbf{x}_{11}, ..., \mathbf{x}_{1n}|\mathbf{x}_{21}, ..., \mathbf{x}_{2n}|......|\mathbf{x}_{L1}, ..., \mathbf{x}_{Ln}].$$

We consider an observation vector $\mathbf{y} \in \mathbb{R}^N$ of unknown class as a linear combination of the training vectors as

$$\mathbf{y} = \sum_{i=1}^{L} \sum_{j=1}^{n} \alpha_{ij} \mathbf{x}_{ij} \qquad (1)$$

with coefficients $\alpha_{ij} \in \mathbb{R}$. The above equation can be written more compactly as

$$\mathbf{y} = \mathbf{D}\alpha, \qquad (2)$$

where

$$\alpha = [\alpha_{11}, ..., \alpha_{1n}|\alpha_{21}, ..., \alpha_{2n}|......|\alpha_{L1}, ..., \alpha_{Ln}]^T \qquad (3)$$

and $.^T$ denotes the transposition operation. We assume that given sufficient training samples of the k^{th} class, $\mathbf{D}_k$, any new test image $\mathbf{y} \in \mathbb{R}^N$ that belongs to the same class will lie approximately in the linear span of the training samples from the class k. This implies that most of the coefficients not associated with class k in (3) will be close to zero. Hence, α is be a sparse vector.

B. Sparse Recovery

In order to represent an observed vector $\mathbf{y} \in \mathbb{R}^N$ as a sparse vector α, one needs to solve the system of linear equations (2). Typically $L.n \gg N$ and hence the system of linear equations (2) is under-determined and has no unique solution. It has been shown that if α is sparse enough and $\mathbf{D}$ satisfies certain properties, then the sparsest α can be recovered by solving the following optimization problem [30] [31] [32] [33]

$$\hat{\alpha} = \arg \min_{\alpha'} \| \alpha' \|_1 \quad \text{subject to} \quad \mathbf{y} = \mathbf{D}\alpha', \qquad (4)$$

where $\| x \|_1 = \sum_i |(x_i)|$. This problem is often known as Basis Pursuit (BP) and can be solved in polynomial time [34][1]. When noisy observations are given, Basis Pursuit DeNoising (BPDN) can be used to approximate α

$$\hat{\alpha} = \arg \min_{\alpha'} \| \alpha' \|_1 \quad \text{subject to} \quad \|\mathbf{y} - \mathbf{D}\alpha'\|_2 \leq \varepsilon, \qquad (5)$$

where we have assumed that the observations are of the following form

$$\mathbf{y} = \mathbf{D}\alpha + \eta \qquad (6)$$

with $\| \eta \|_2 \leq \varepsilon$.

Finally, in certain cases, greedy algorithms such as orthogonal matching pursuit [35] can also be used to recover sparse images. Such an algorithm for face recognition has been proposed in [36].

C. Sparse Recognition

Given an observation vector $\mathbf{y}$ from one of the L classes in the training set, we compute its coefficients $\hat{\alpha}$ by solving either (4) or (5). We perform classification based on the fact that high values of the coefficients $\hat{\alpha}$ will be associated with the columns of $\mathbf{D}$ from a single class. We do this by comparing how well the different parts of the estimated coefficients, $\hat{\alpha}$, represent $\mathbf{y}$. The minimum of the representation error or the residual error is then used to identify the correct class. The residual error of class k is calculated by keeping the coefficients associated with that class and setting the coefficients not associated with class k to zero. This can be done by introducing a

---

[1]Note that the ℓ_1 norm is an approximation of the the ℓ_0 norm. The approximation is necessary because the optimization problem in (4) with the ℓ_0 norm (which seeks the sparsest α) is NP-hard and computationally difficult to solve.

characteristic function, $\Pi_k : \mathbb{R}^n \to \mathbb{R}^n$, that selects the coefficients associated with the k^{th} class as follows

$$r_k(\mathbf{y}) = \|\mathbf{y} - \mathbf{D}\Pi_k(\hat{\alpha})\|_2. \qquad (7)$$

Here the vector Π_k has value one at locations corresponding to the class k and zero for other entries. The class, d, which is associated with an observed vector, is then declared as the one that produces the smallest approximation error

$$d = \arg \min_k r_k(\mathbf{y}). \qquad (8)$$

We now summarize the sparse recognition algorithm as follows:
Given a matrix of training samples $\mathbf{D} \in \mathbb{R}^{N \times (n.L)}$ for L classes and a test sample $\mathbf{y} \in \mathbb{R}^N$:

1) Solve the BP (4) or BPDN (5) problem.
2) Compute the residual using (7).
3) Identify $\mathbf{y}$ using (8).

D. Image quality measure

For classification, it is important to be able to detect and then reject the test samples of poor quality. To decide whether a given test sample has good quality, we use the notion of Sparsity Concentration Index (SCI) proposed in [4]. The SCI of a coefficient vector $\alpha \in \mathbb{R}^{(L.n)}$ is defined as

$$SCI(\alpha) = \frac{\frac{L. \max \|\Pi_i(\alpha)\|_1}{\|\alpha\|_1} - 1}{L - 1}. \qquad (9)$$

SCI takes values between 0 and 1. SCI values close to 1 correspond to the case where the test image can be approximately represented by using only images from a single class. The test vector has enough discriminating features of its class, so has high quality. If SCI = 0 then the coefficients are spread evenly across all classes. So the test vector is not similar to any of the classes and has of poor quality. A threshold can be chosen to reject the iris images with poor quality. For instance, a test image can be rejected if $SCI(\hat{\alpha}) < \lambda$ and otherwise accepted as valid, where λ is some chosen threshold between 0 and 1.

IV. BAYESIAN FUSION BASED IMAGE SELECTION AND RECOGNITION

Different regions of the iris have different qualities [11]. So instead of recognizing the entire iris image directly, we recognize the different regions separately and combine the results depending on the quality of the region. This reduces the computational complexity of the above method as the size of the dictionary is greatly reduced, and the recognition of the different regions can be done in parallel. Also, since occlusions affect only local regions on the iris which can only lower the quality of certain regions, the robustness of the recognition algorithm to occlusion due to eyelids and eye lashes is improved. A direct way of doing this would be to recognize the sectors separately and combine the results by voting [37]. This, however, does not account for the fact that different regions are recognized with different confidences. In what follows, we propose a score level fusion approach for recognition where we combine the recognition results of

different sectors based on the recognition confidence using the corresponding SCI values. Fig. 2 illustrates the different steps involved in the proposed approach.
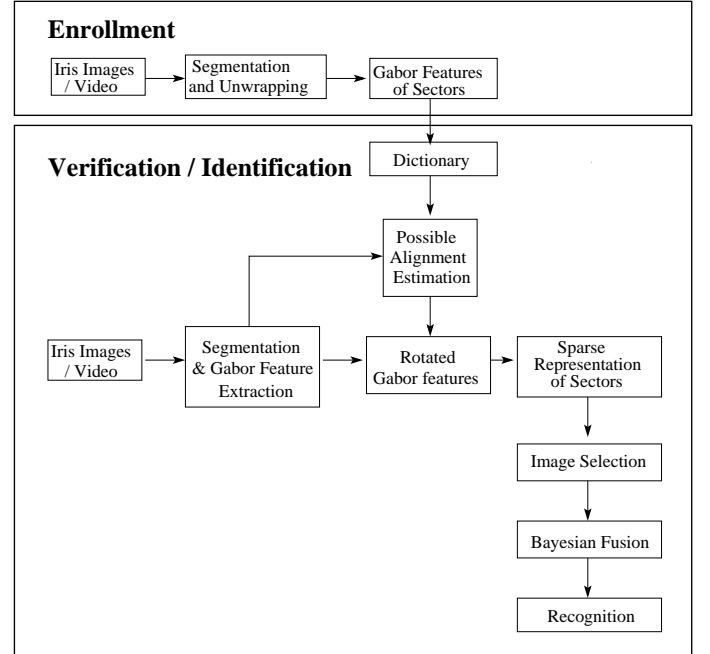

Fig. 2. A block diagram illustrating the Bayesian Fusion based image selection and recognition.

Consider the iris recognition problem with L distinct classes. Let $\mathbf{C} = \{c_1, c_2, \ldots, c_L\}$ be the class labels. Let $\mathbf{y}$ be the test vector whose identity is to be determined. Let us divide the vector $\mathbf{y}$ into $\hat{M}$ non-overlapping regions, each called a sector. Each of the sectors is individually solved using the sparse representation-based recognition algorithm discussed in section III. The sectors with SCI values below the threshold are rejected. Let M be the number of sectors retained, where $M \leq \hat{M}$. Let $d_1, d_2, \ldots, d_M$ be the class labels of the retained sectors. Ideally, if the data is noise free, all the returned labels will be equal to the true label c. That is,

$$d_1 = d_2 = \ldots = d_M = c.$$

However, in the presence of noise in the training and test iris images, the returned labels will not necessarily be the same. Let $\mathbb{P}(d_i|c)$ be the probability of the i^{th} sector returns a label d_i when the true class is c. It is reasonable to assume that the probability of the recognition system returning the true label c is high. But given the noise in the iris images, all the classes other than c will still have a low probability of being identified as the true class. SCI is a measure of the confidence in recognition, so the higher the SCI value, the higher the probability that the true class will be the same as the class suggested by the recognition system. So a reasonable model for the likelihood is

$$\mathbb{P}(d_i|c) = \begin{cases} \frac{t_1^{SCI(d_i)}}{t_1^{SCI(d_i)} + (L-1).t_2^{SCI(d_i)}} & if \ d_i = c, \\ \frac{t_2^{SCI(d_i)}}{t_1^{SCI(d_i)} + (L-1).t_2^{SCI(d_i)}} & if \ d_i \neq c \end{cases} \qquad (10)$$

where t_1 and t_2 are positive constants such that

$$t_1 > t_2 > 1$$

The numerator gives a higher probability value to the correct class, and the denominator is a normalizing constant. The condition (IV) ensures that the probability of the true class increases monotonically with the SCI value of the sector. Thus, this likelihood function satisfies the two constraints mentioned above.

The maximum aposteriori estimate (MAP) of the class label given the noisy individual sector labels is given by

$$\tilde{c} = \arg\max_{c \in \mathbf{C}} \mathbb{P}(c|d_1, d_2, \ldots, d_M) \qquad (11)$$

Assuming the prior probabilities of the classes are uniform, we obtain

$$\tilde{c} = \arg\max_{c \in \mathbf{C}} \mathbb{P}(d_1, d_2, \ldots, d_M|c)$$

Conditioned on the true class, the uncertainty in the class labels is only due to the noise in the different sectors, which are assumed to be independent of each other. So

$$\tilde{c} = \arg\max_{c \in \mathbf{C}} \prod_{j=1}^{M} \mathbb{P}(d_j|c)$$
$$= \arg\max_{c \in \mathbf{C}} t_1^{\sum_{j=1}^{M} SCI(d_j).\delta(d_j=c)} . t_2^{\sum_{j=1}^{M} SCI(d_j).\delta(d_j \neq c)} \qquad (12)$$

where $\delta(.)$ is the Kronecker delta function. Since $t_1 > t_2$, the solution to (12) is same as

$$\tilde{c} = \arg\max_{c \in \mathbf{C}} \sum_{j=1}^{M} SCI(d_j).\delta(d_j = c) \qquad (13)$$

Let us define the Cumulative SCI (CSCI) of a class c_l as

$$CSCI(c_l) = \frac{\sum_{j=1}^{M} SCI(d_j).\delta(d_j = c_l)}{\sum_{j=1}^{M} SCI(d_j)} \qquad (14)$$

So

$$\tilde{c} = \arg\max_{c \in \mathbf{C}} CSCI(c) \qquad (15)$$

CSCI of a class is the sum of the SCI values of all the sectors identified by the classifier as belonging to that class. Therefore, the optimal estimate is the class having the highest CSCI.

V. Iris Recognition from video

In this section, we illustrate how our method can be extended to perform recognition from iris videos. Let $Y = \{\mathbf{y}^1, \mathbf{y}^2, \ldots, \mathbf{y}^J\}$ be the J vectorized frames in the test video. As before, each frame is divided into $\hat{M}$ sectors and recognized separately by the sparse recognition algorithm. Let M_i be the number of sectors retained by the selection scheme in the i^{th} frame. Let $\mathbf{y}_j^i$ be the j^{th} retained sector in the i^{th} frame. Using a derivation similar to the one given in Section IV, we can derive the MAP estimate as

$$\tilde{c} = \arg\max_{c \in \mathbf{C}} \sum_{i=1}^{J} \sum_{j=1}^{M_i} SCI(d_j^i).\delta(c = d_j^i) \qquad (16)$$

where d_j^i is the class label assigned by the classifier to $\mathbf{y}_j^i$. (16) can be alternatively written as

$$\tilde{c} = \arg\max_{c \in \mathbf{C}} CSCI(c) \qquad (17)$$

where CSCI of a class c_l is given by

$$CSCI(c_l) = \frac{\sum_{i=1}^{J} \sum_{j=1}^{M_i} SCI(d_j^i).\delta(d_j^i = c_l)}{\sum_{i=1}^{J} \sum_{j=1}^{M_i} SCI(d_j^i)}. \qquad (18)$$

As before, the MAP estimate consists of selecting the class having the highest cumulative SCI value, with the difference that the sectors of all the frames in the test video will be used while computing the CSCI of each class. Note that unlike existing feature level and score level fusion methods available for iris recognition, the CSCI incorporates the quality of the frames into the matching score. Hence, when the frames in the video suffer from acquisition artifacts like blurring, occlusion and segmentation errors, the proposed matching score gives higher weights to the good frames, at the same time, suppressing the evidence from the poorly acquired regions in the video.

The different modes of operation of the proposed algorithm are illustrated in Fig. 3. Both the probe and the gallery can be separate iris images or iris videos. The iris images are segmented and unwrapped to form rectangular images. The Gabor features of the different sectors are computed, and sparse representation-based recognition algorithm described in section III-C is used to select the good iris images. The good sectors are separately recognized and combined to obtain the class of probe image or video as described above.

VI. Handling Alignment

Due to rotation of the head with respect to the camera, the captured test iris image may be rotated with respect to the training images. To obtain a good recognition performance, it is important to align the test images before recognition. In this section, we propose a two stage approach for iris feature alignment. In the first stage, we estimate the best K alignments for each test vector using matched filters and then obtain an alignment invariant score function, based on the Bayesian fusion framework introduced above.

A. Matched Filter Based Alignment Estimation

Let $\mathbf{y}$ be the test vector to be recognized. Let $\hat{A}$ be the number of possible alignments of the test vector. A matched filter is designed for each alignment, whose impulse response is equal to the corresponding shifted version of $\mathbf{y}$. Let $\mathbf{h}_i$ be the impulse response of the i^{th} matched filter, and $\mathbf{H}$ be the set of all possible impulse responses.

$$\mathbf{H} = \{\mathbf{h}_1, \mathbf{h}_2, \ldots, \mathbf{h}_{\hat{A}}\} \qquad (19)$$

Let e_{ijk} be the sum of squared error between i^{th} matched filter impulse response and j^{th} training image of the k^{th} class.

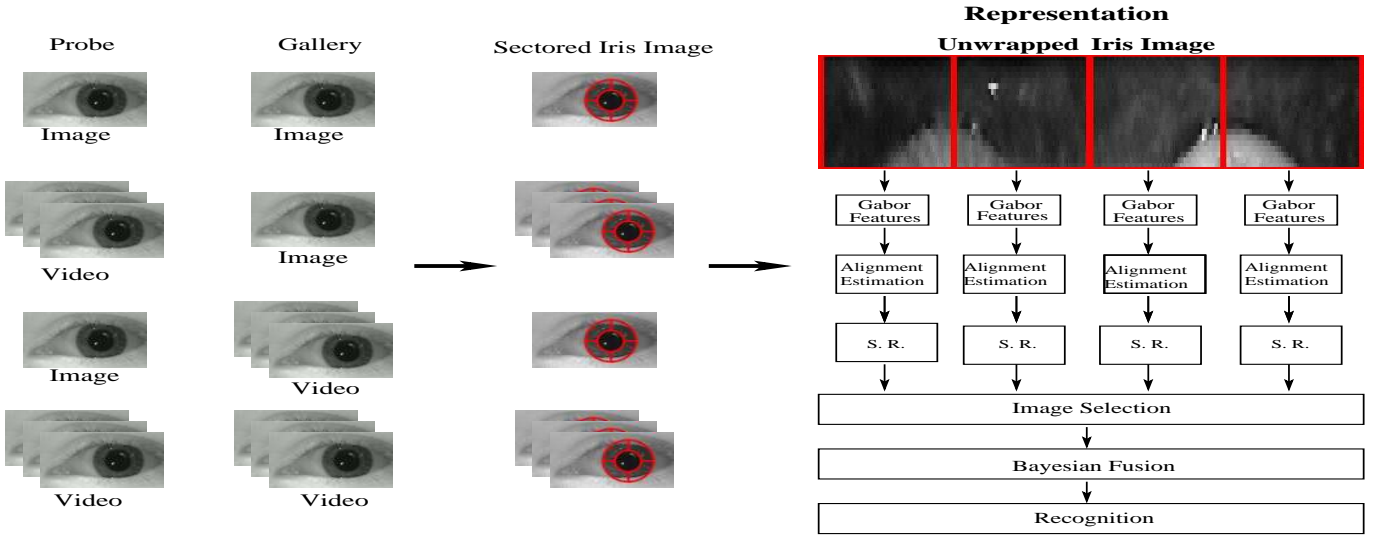$$e_{ijk} = \|\mathbf{h}_i - \mathbf{x}_{kj}\|_2^2 \qquad (20)$$

Fig. 3. A block diagram illustrating the different modes of operation of the proposed algorithm. Both the probe and the gallery can be individual iris images or iris video. Here, S.R. stands for Sparse Representation.

The alignment error associated with the i^{th} alignment is computed as

$$e_i = \min_{k=1,2,\ldots L, j=1,2,\ldots n} e_{ijk} \qquad (21)$$

The top K alignments are selected as the ones producing the least alignment error e_i.

B. Score Estimation Robust to Alignment Errors

From each test vector $\mathbf{y}$, we can generate K new test vectors by shifting it according to the corresponding alignments obtained from the method described above. So instead of the J original frames in the video, we now have JK frames. Using arguments similar to the ones in the previous section, we can obtain the CSCI of the l^{th} class c_l as

$$CSCI(c_l) = \frac{\sum_{i=1}^{JK} \sum_{j=1}^{M_i} SCI(d_j^i) . \delta(d_j^i = c_l)}{\sum_{i=1}^{JK} \sum_{j=1}^{M_i} SCI(d_j^i)}. \qquad (22)$$

where M_i are the number of sectors retained in the i^{th} frame. The MAP estimate of the output class is the one with the highest CSCI value. Note that this score estimation handles the global alignment errors and not the local deformations in the iris pattern. Since our method weighs different sectors based on their quality, sectors having significant local deformations will not have high influence on the final CSCI value due to their lower quality.

VII. SECURE IRIS BIOMETRIC

For a biometric system to be deployed successfully in a practical application, ensuring security and privacy of the users is essential. In this section, we propose two cancelable methods to improve security of our recognition system.

A. Cancelability Through Random Projections

The idea of using Random Projections (RP) for cancelability in biometrics has been previously introduced in [28], [38], [39]. In [28] and [38], RPs of discriminative features were used for cancelability in face biometrics. RPs on different regions of the iris were applied for cancelability in [39]. In what follows, we show how RPs can be extended into the sparse representation-based approach for ensuring cancelability.

Let Φ be an $m \times N$ random matrix with $m \leq N$ such that each entry $\phi_{i,j}$ of Φ is an independent realization of q, where q is a random variable on a probability measure space (Ω, ρ). Consider the following observations:

$$\mathbf{a} \doteq \Phi \mathbf{y} = \Phi \mathbf{D} \alpha + \eta', \qquad (23)$$

where $\eta' = \Phi \eta$ with $\| \eta' \|_2 \leq \varepsilon'$. $\mathbf{a}$ can be thought of as a transformed version of the biometric $\mathbf{y}$. One must recover the coefficients α to apply the sparse recognition method explained in section III. As m is smaller than N, the system of equations (23) is underdetermined and a unique solution of α is not available. Given the sparsity of α, one can approximate α by solving the BPDN problem. It has been shown that for sufficiently sparse α and under certain conditions on $\Phi \mathbf{D}$, the solution to the following optimization problem will approximate the sparsest near-solution of (23) [40]

$$\hat{\alpha} = \arg \min_{\alpha'} \| \alpha' \|_1 \quad \text{s. t.} \quad \|\mathbf{a} - \Phi \mathbf{D} \alpha'\|_2 \leq \varepsilon'. \qquad (24)$$

One sufficient condition for (24) to stably approximate the sparsest solution of (23), is the Restricted Isometry Property (RIP)[41], [32], [33]. A matrix $\Phi \mathbf{D}$ satisfies the RIP of order K with constants $\delta_K \in (0, 1)$ if

$$(1 - \delta_K) \| v \|_2^2 \leq \| \Phi \mathbf{D} v \|_2^2 \leq (1 + \delta_K) \| v \|_2^2 \qquad (25)$$

for any v such that $\| v \|_0 \leq K$. When RIP holds, ΦD approximately preserves the Euclidean length of K-sparse vectors. When D is a deterministic dictionary and Φ is a random matrix, the following theorem on the RIP of ΦD can be stated.

Theorem 1. *([40]) Let $D \in \mathbb{R}^{N \times (n.L)}$ be a deterministic dictionary with restricted isometry constant $\delta_K(D), K \in \mathbb{N}$. Let $\Phi \in \mathbb{R}^{m \times N}$ be a random matrix satisfying*

$$P\left(|\|\Phi v\|^2 - \|v\|^2| \geq \varsigma \|v\|^2\right) \leq 2e^{-c\frac{n}{2}\varsigma^2}, \quad \varsigma \in (0, \frac{1}{3}) \quad (26)$$

for all $v \in \mathbb{R}^N$ and some constant $c > 0$ and assume

$$m \geq C\delta^{-2}\left(K \log((n.L)/K) + \log(2e(1 + 12/\delta)) + t\right) \quad (27)$$

for some $\delta \in (0, 1)$ and $t > 0$. Then, with probability at least $1 - e^{-t}$, the matrix ΦD has restricted isometry constant

$$\delta_K(\Phi D) \leq \delta_K(D) + \delta(1 + \delta_K(D)). \quad (28)$$

The constant satisfies $C \leq 9/c$.

The above theorem establishes how the isometry constants of D are affected by multiplication with a random matrix Φ. Note that one still needs to check the isometry constants for the dictionary D to use this result. However, for a given dictionary, D, it is difficult to prove that D satisfies a RIP. One can alleviate this problem by using the phase transition diagrams [42], [43]. See section VII-A for more details.

The following are some matrices that satisfy (26) and hence can be used as random projections for cancelability.

- $m \times N$ random matrices Φ whose entries $\phi_{i,j}$ are independent realizations of Gaussian random variables $\phi_{i,j} \sim N\left(0, \frac{1}{m}\right)$.
- Independent realizations of ± 1 Bernoulli random variables

$$\phi_{i,j} \doteq \begin{cases} +1/\sqrt{m}, & \text{with probability } \frac{1}{2} \\ -1/\sqrt{m}, & \text{with probability } \frac{1}{2}. \end{cases}$$

- Independent realizations of related distributions such as

$$\phi_{i,j} \doteq \begin{cases} +\sqrt{3/m}, & \text{with probability } \frac{1}{6} \\ 0, & \text{with probability } \frac{2}{3} \\ -\sqrt{3/m}, & \text{with probability } \frac{1}{6}. \end{cases}$$

- Multiplication of any $m \times N$ random matrix Φ with a deterministic orthogonal $N \times N$ matrix $\tilde{D}$, i.e. $\Phi\tilde{D}$.

Note that RPs meet the various constraints required for cancelability, mentioned in Section I. By using different RP matrices, we can issue different templates for different applications. If a transformed pattern is compromised, we can reissue a new pattern by applying a new random projection to the iris vector. The RIP properties together with the sparsity of α ensure that the recognition performance is preserved. In the application database, only the transformed dictionary ΦD is stored. If a hacker illegally obtains the transformed dictionary ΦD and the transformed iris patterns of the user, a, he or she will have access to the person's identity. However, it is extremely difficult to obtain the matrix D from ΦD, and without D one cannot obtain the original iris patterns
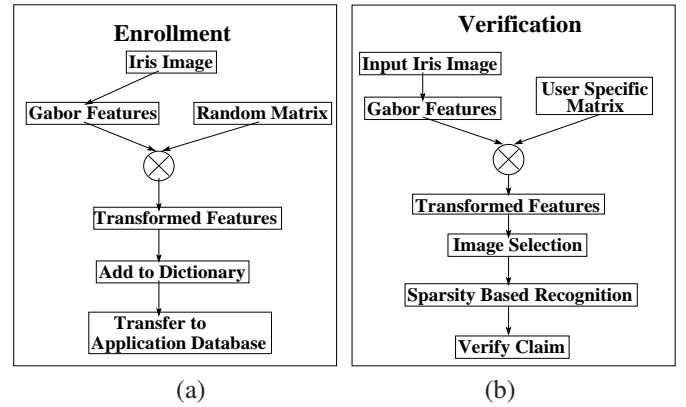

Fig. 4. Block Diagram of the Random Projections based cancelable system.

y. Hence, our cancelable scheme is non-invertible as it is not possible to obtain the original iris patterns from the transformed patterns. Furthermore, since our method is based on pseudo-random number generation, we only consider the state space corresponding to the value taken by the seed of the random number generator. Hence, instead of storing the entire matrix, one only needs to store the seed used to generate the RP matrix.

B. Cancelability through Random Permutations of dictionary columns

As explained in section III, when the iris image has good quality, only the training images corresponding to the correct class will have high coefficients. If the training images of different classes are randomly arranged as columns of the dictionary, both the dictionary and the order of the training images are required for correct recognition. In this section, we explain how this idea can enhance the security of our iris recognition system.

When a new user is enrolled, his training images are divided into sectors and placed at random locations in the dictionary. In Fig. 5, we show the dictionary for a trivial example of four users. Note that the different sectors of each training image of the user are kept at different random locations in the dictionary. Without prior knowledge of these locations, it is impossible to perform recognition.

An array indicating the column numbers of the training images of the correct class is generated for each user. This array is stored in a hash table, and the corresponding hash code is given to the user during enrollment. During verification, the system acquires the iris image of the person and extracts the features. For each sector of the iris vector, the sparse coefficients are obtained using this shuffled dictionary, as explained in section III. The user also has to present the hash code to the system. Using the hash code, the indices of training images are obtained from the hash table and the coefficients belonging to different classes are grouped. Then, SCI is computed and used to retain or reject the images. If the image is retained, the CSCI values of the different classes are computed and the class having the lowest CSCI value is assigned as the class label of the user, as explained
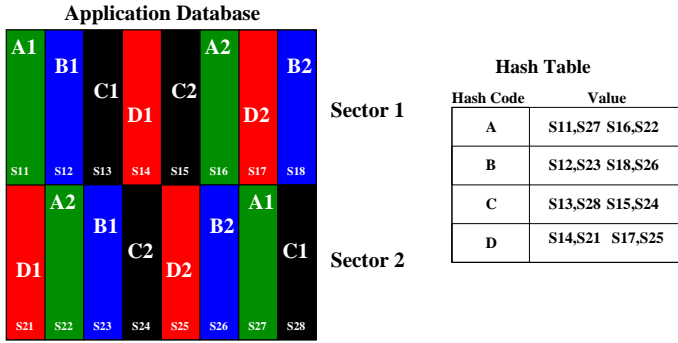
Application Database

Fig. 5. Sample Dictionary and hash table for a four user example. The four users A, B, C and D are indicated by colors green, blue, black and red, respectively. A1 and A2 are the two training images corresponding to the first user. Sij denote that the j^{th} location and the i^{th} sector. $D1$ at $S14$ means that the first sector of the user D is at location $S14$.

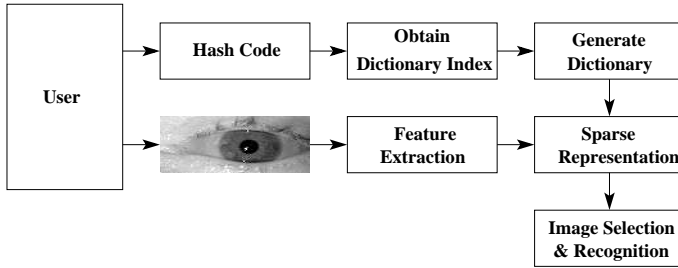in section IV. A block diagram of the security scheme is presented in Fig. 6

Fig. 6. Block Diagram of the proposed cancelability scheme using random permutations.

If the hash code presented is incorrect, then the obtained indices of the training images for each class will be wrong. So the coefficients will be grouped in a wrong way, and all the classes will have similar energy leading to a low SCI value and the subsequent rejection of the image. Even if by chance, one of the classes happened to have high energy and the image is retained, the probability of that class being the correct class is very low ($\frac{1}{N}$). Thus, with high probability, the user will not be verified. Hence, if a hacker illegally acquires the iris patterns of a legitimate user, without having the hash code, he or she will not be able to access the system. Also, even if the hacker obtains the iris dictionary stored in the application database, the iris patterns of the user cannot be accessed without knowing the correct hash codes, because different sectors of an iris patterns reside at different random locations. If the hash code is compromised, the dictionary indices of the user can then be stored at a new location, and a new hash code can be issued to the user. Also, different applications can have different dictionaries. Thus, the user will have a different hash code for each application, preventing cross matching.

It should be noted that the additional security and privacy introduced by these techniques come at the expense of storing additional seed values. In applications requiring higher security, this can be stored with the user, so that a hacker will not

get the original templates even if he gets hold of the cancelable patterns in the template database. For applications with greater emphasis on usability, the seed can be stored securely in the template database, so that the user will not have to carry it.

VIII. RESULTS AND DISCUSSION

In the following subsections, we present iris image selection, recognition and cancelability results on the ICE2005 dataset [3], ND-IRIS-0405 (ND) dataset [44] and the MBGC videos [45]. The ND dataset is a superset of the ICE2005 and ICE2006 iris datasets. It contains about sixty five thousand iris images belonging to three hundred and fifty six persons, with a wide variety of distortions, facilitating the testing and performance evaluation of our algorithm. In all of our experiments, we employed a highly efficient algorithm suitable for large scale applications, known as the Spectral Projected Gradient (SPGL1) algorithm [46], to solve the BP and BPDN problems.

A. Empirical verification of ℓ_0/ℓ_1 equivalence

Our sparse recognition algorithm's performance depends on certain conditions on the dictionary such as incoherence and RIP. However, as stated earlier, it is very difficult to prove any general claim that $\mathbf{D}$, $\mathbf{GD}$, $\Phi\mathbf{D}$, or $\Phi\mathbf{GD}$ satisfies a RIP or an incoherence property. To address this, one can use the phase transition diagrams [42]. A phase transition diagram provides a way of checking ℓ_0/ℓ_1 equivalence, indicating how sparsity and indeterminacy affect the success of ℓ_1 minimization [42], [43].
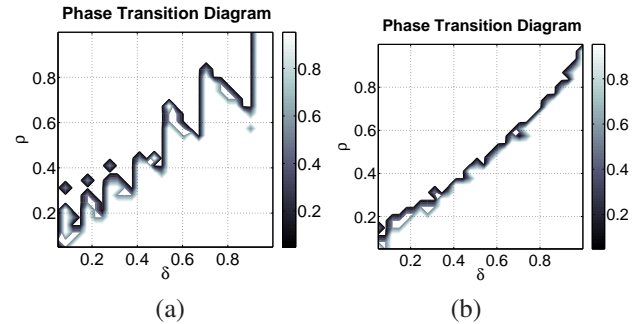

Fig. 7. Phase transition diagrams corresponding to the case when the dictionary is (a) $\mathbf{GD}$ and (b) $\Phi\mathbf{GD}$, where $\mathbf{G}$ is the Gabor transformation matrix and Φ is the random projection matrix for cancelability. In both figures, we observe a phase transition from lower region where the ℓ_0/ℓ_1 equivalence holds, to the upper region, where one must use combinatorial search to recover the sparsest solution.

Let $\delta = \frac{M}{N}$ be a measure of undersampling factor, and $\rho = \frac{K}{M}$ be a measure of sparsity. A plot of the pairing of the variables δ and ρ describes a two-dimensional phase space $(\delta, \rho) \in [0, 1]^2$. The values of δ and ρ ranged through 40 equispaced points in the interval $[0, 1]$ and $N = 800$. At each point on the grid, we recorded the mean number of coordinates at which original and reconstruction differed by more than 10^{-3}, averaged over 20 independent realizations (see [42], [43] for more details).

In Fig. 7 (a) and (b), we show the phase transition diagrams corresponding to the case when the dictionary is $\mathbf{GD}$ and

$\Phi\mathbf{GD}$, respectively. Here, $\mathbf{G}$ is the Gabor transformation matrix and Φ is an $m \times N$ matrix whose entries $\phi_{i,j}$ are independent realizations of Gaussian random variables $\phi_{i,j} \sim N\left(0, \frac{1}{m}\right)$. For each value of δ, the values of ρ below the curve, are the ones where the ℓ_0/ℓ_1 equivalence holds. As can be observed, for most values of δ, there is atleast one value of ρ below the curve, satisfying the equivalence. So the vector α can be recovered if it is sparse enough and enough measurements are taken.

B. Image Selection and Recognition

In this section, we evaluate our selection and recognition algorithms on ND and ICE2005 datasets. To illustrate the robustness of our algorithm to occlusion due to eyelids and eyelashes, we perform only a simple iris segmentation scheme, detecting just the pupil and iris boundaries and not the eyelids and eye lashes. We use the publicly available code of Masek *et al.* [47] for detecting these boundaries.

1) Variation of SCI with common distortions during image acquisition: To study the variation of SCI in the presence of common distortions during image acquisition like occlusion and blur, we simulate them on the clean iris images from the ND dataset.

Description of the Experiment: We selected fifteen clean iris images of the left eye of eighty persons. Twelve such images per person formed the gallery and distortions were simulated on the remaining images to form the probes. We consider seven different levels of distortion for each case, with level one indicating no distortion and level seven indicating maximum distortion. We obtain the dictionary using the gallery images, and evaluate the SCI of the various sectors of the test images.

Fig. 8 shows some of the simulated images from the ND dataset. The first column includes images with distortion level one (no distortion). The middle column contains images with distortion level three (moderate distortions). The right most column contain images with distortion level five (high distortion). The first row contains images with blur while the second contains images with occlusion. Images with simulated segmentation error and specular reflections are shown in the third and fourth rows respectively.

Fig. 9 (a) illustrates the variation of SCI with the common acquisition distortions. It can be observed that good images have high SCI values whereas the ones with distortion have lower SCI values. So by suitably thresholding the SCI value of the test image, we can remove the bad images before the recognition stage. The relative stability in SCI values with occlusion and specular reflection demonstrates the increased robustness attained by our algorithm, by separately recognizing the individual sectors and combining the results, as mentioned in section IV.

2) Image Selection results on the ND dataset: In this section, we illustrate the performance of our image selection algorithm on images from the ND dataset.

Description of the Experiment: We selected the left iris images of eighty subjects that had sufficiently large number of iris images with different distortions like blur, occlusion
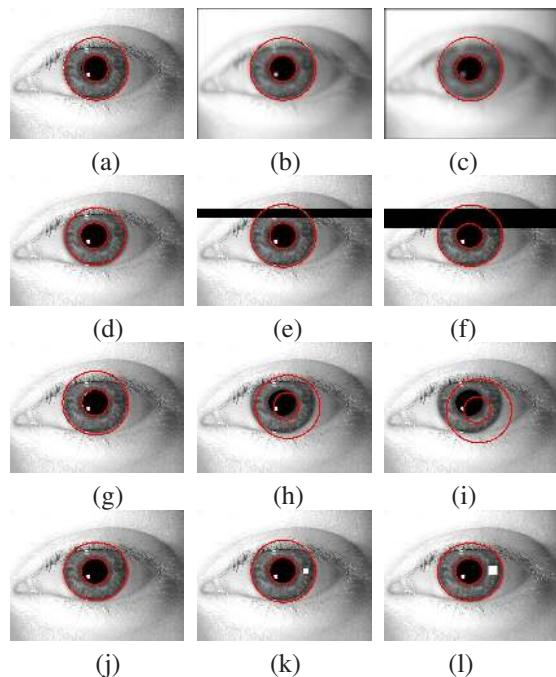

Fig. 8. Simulated Distortions on the images from the ND dataset. The detected pupil and iris boundaries are indicated as red circles.

and segmentation errors. Fifteen clean images per person were hand chosen to form the gallery. Up to fifteen images with blur, occlusion and segmentation errors were also selected. As mentioned before, we perform a simple segmentation scheme, retaining possible occlusion due to eyelids and eyelashes in the iris vector. The Gabor features of the iris vector form the input. Our algorithm creates the dictionary, finds the sparse representation for each test vector, evaluates the SCI of the sectors, and rejects the images for which all the sectors have SCI value below a threshold of 0.6.

Measure the selection performance : The quality of the input iris feature vector should be a function of the performance of the recognition algorithm on that sample [1]. An ideal image selection algorithm should retain images, which can be correctly recognized by the recognition algorithm, and reject the ones on which the subsequent recognition algorithm will perform poorly. To measure it, we define the Modified False Positive Rate (MFR) and a Modified Verification Rate (MVR) as follows. Modified False Positive rate is the fraction of the test vectors retained by the image selection algorithm, which are wrongly classified by the subsequent recognition algorithm. Modified Verification Rate is defined as the fraction of the images correctly classified by the recognition algorithm, which are retained by the selection scheme. To obtain these values, we find the CSCI for each test sample and also the class assigned to the samples by our algorithm. We obtain the Receiver Operating Characteristics (ROC) of the image selection algorithm by plotting MVR versus MFR for different values of threshold. Note that this measures the performance of the quality estimation stage and is different from the ROC curve of the recognition algorithm.

$$MFR = \frac{\text{No of Images selected and wrongly classified}}{\text{No of images selected}}$$
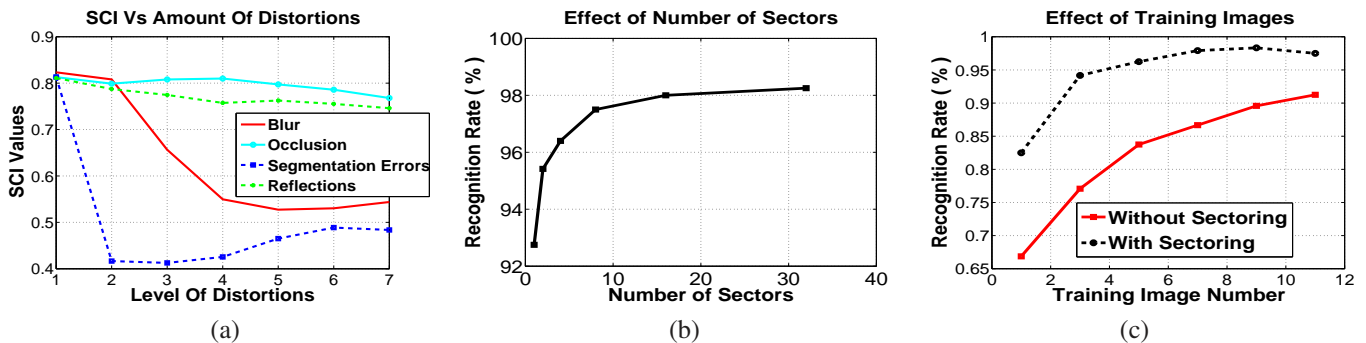
Fig. 9. (a) Plot of the variation in SCI values with common distortions in iris image acquisition. Note that the SCI falls monotonically with increasing levels of blur and segmentation errors in the iris images. It is also robust to occlusions and specular reflections. (b) Plot of the recognition rate versus the number of sectors. Observe the significant improvement in the results as the number of sectors is improved from one to eight. (c) Plot of the recognition rate versus the number of training images. Note that the recognition rate increases monotonically with the number of training images. Also, sectoring achieves the same recognition rate as the case without sectoring using far fewer training images.

$$MVR = \frac{\text{No of Images selected and correctly classified}}{\text{No of images correctly classified}}$$

Fig. 10(a) shows the ROC of our image selection algorithm (black), compared to that using directly the Hamming distance based on the publicly available iris recognition system of Masek *et al.* [47] (red), when the probe images are blurred. Since the data has occlusion, direct application of Masek's algorithm performed poorly. For a fair comparison, we modified the algorithm, recognizing the different sectors of the iris separately and fusing the results through voting. Note that our ROC curve is significantly sharper than that of the Masek's recognition system indicating superior performance.

The effects of occlusion in iris images due to eyelids, eye lashes and specular reflections are illustrated in Fig. 10(b). Images with occlusion were obtained for each of the eighty classes under consideration and used as probes. The ROC curve of our algorithm is shown in black and that of Masek's system appears in red. Note that for the same MFR, the proposed image selection scheme has a higher MVR. This indicates that the proposed selection method retains more images that will be correctly classified by the subsequent recognition algorithm and rejects more images that will be wrongly classified by the recognition algorithm.

To study the effects of segmentation error, the gallery images were verified to be well segmented. Up to fifteen images with segmentation errors were chosen for each person under consideration, which formed the probes. Fig. 10(c) shows the ROC curves of our method (black) and the Masek's one (red) in case of wrongly segmented images. Again, using our image selection algorithm improves the performance of the system even with wrongly segmented images, a feature lacking in many existing quality estimation methods.

3) Recognition Results on images from the ND dataset: In this section, we illustrate the performance of our recognition algorithm on images from the ND dataset.

Performance on clean images - Description of the Experiment: Eighty subjects were selected from the dataset. Fifteen clean images of the left iris were hand selected for each person. Of these fifteen images per person, twelve were randomly selected to form the gallery and the remaining three images per person were used as probes. No image selection

is performed because we want to evaluate the performance of the recognition algorithm separately.

We compare our algorithm to a nearest neighbor based recognition algorithm (NN) that uses the Gabor features and the Masek's implementation. Since we use tough segmentation conditions retaining the eyelids and eye lashes in the iris vector, direct application of NN and Masek's method produced poor results. For a fair comparison, we divided the iris images into different sectors, obtained the results using these methods separately on each sectors and combined the results by voting. We obtained a recognition rate of 99.15% when compared to 98.33% for the NN and 97.5% for the Masek's method.

Performance on poorly acquired images - Description of the Experiment - To evaluate the recognition performance of our algorithm on poorly acquired images, we hand picked images with blur, occlusion and segmentation errors as explained in the previous section. Fifteen clean images per person were used to form the gallery. Probes containing each type of distortion were applied separately to the algorithm. We perform image selection followed by recognition. The recognition rates are reported in Table. II.

TABLE I
RECOGNITION RATE ON ND DATASET

Image Quality	NN	Masek's Implementation	Proposed Method
Good	98.33	97.5	99.15
Blurred	95.42	96.01	98.18
Occluded	85.03	89.54	90.44
Seg. Error	78.57	82.09	87.63

In Fig. 11, we display the iris images having the least SCI value for the blur, occlusion and segmentation error experiments performed on the real iris images in the ND dataset as mentioned above. As can be observed, images with low SCI values suffer from high amounts of distortion.

4) Recognition Performance on the ICE 2005 Dataset: In this section, we compare the performance of our algorithm with the existing results on the ICE 2005 dataset corresponding to Experiment 1. Experiment 1 has 1425 iris images corresponding to 120 different classes.
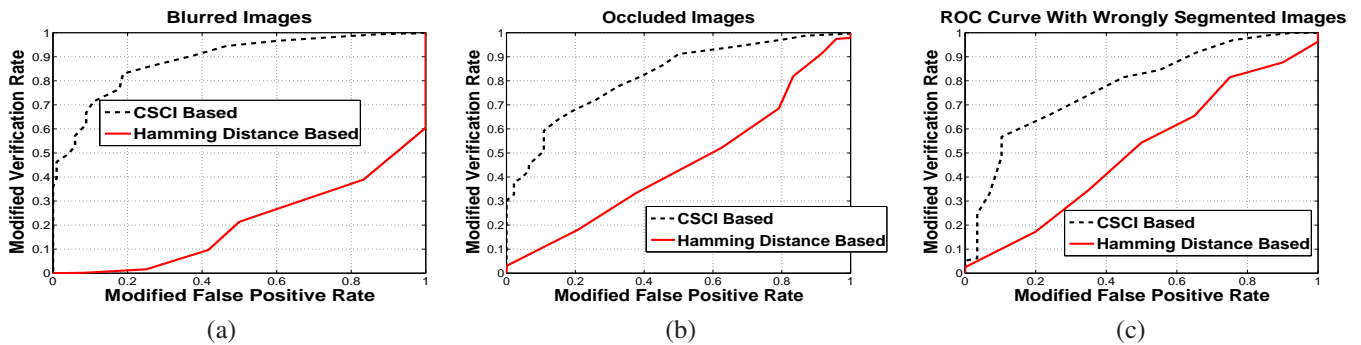
Fig. 10.   Comparison of the ROC curves of the proposed image selection algorithm (CSCI Based) and one using Hamming distance as the quality measure(Hamming Distance Based) using clean iris images in the gallery and probe images containing (a) Blurring (b) Occlusions and (c) Segmentation Errors. Note that CSCI based image selection performs significantly better than Hamming distance based selection when the image quality is poor.
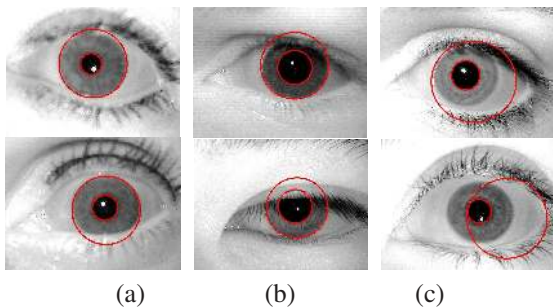

Fig. 11.   Iris images with low SCI values in the ND dataset. Note that the images in (a), (b) and (c) suffer from high amounts of blur, occlusion and segmentation errors respectively .

Description of the Experiment : We have used ten images per class in the gallery and remaining iris images as the test vectors. We perform segmentation using Masek's code and apply the Gabor features of the segmented iris images to our recognition algorithm. No image selection was performed. We compare our performance with existing results in Table II, where the verification rates are indicated at a false acceptance rate of 0.001. The results of the existing methods are obtained from [48].

TABLE II
VERIFICATION RATE AT AN FAR OF 0.001 ON THE ICE 2005 DATASET

Method	Verification Rate (%)
Pelco	96.8
WVU	97.9
CAS 3	97
CAS 1	97.8
CMU	99.5
SAGEM	99.8
Proposed Method	98.13

5) Dependence of recognition rate on the number of sectors: Fig. 9 (b) plots the variation of the recognition rates for the proposed method with changes in the number of sectors. As can be observed, the performance of the recognition system improves significantly as the number of sectors is increased from one to eight. Beyond eight, the recognition rate does not increase significantly.

6) Effect of the number of training images on performance: In this section, we study the effect of the number of training images on recognition rate of our algorithm. We vary the number of training images from one per class to eleven per class on the ND dataset. The test images consisting of three iris images per person are used to test each of these cases. The variation of recognition rate is plotted in Fig. 9 (c) for the case of no sectoring and sectoring with eight sectors respectively. As can be observed, recognition performance increases with the number of training images. This is hardly surprising as our assumption that the training images span the space of testing images becomes more valid as the number of training images increases. In unconstrained iris recognition systems which we are interested in, this is not a bottle neck because we can obtain a significant number of iris images from the incoming iris video. Also, sectoring achieves the same recognition rate as the non-sectoring case with a much lower number of training images.

7) CSCI as a measure of confidence in recognition: We have empirically observed that the higher the CSCI value for the test image, the higher the probability that it is correctly classified. This is illustrated in Fig. 12 (a). This observation is expected as high CSCI means that the reconstructed vector in most of the sectors will be sparse. If the training images span the space of possible testing images, the training images of the correct class will have high coefficients. So the only possible sparse vector is the one in which the correct class has high coefficients and others have zero coefficients, which will be correctly classified by our algorithm.

C. Cancelability Results using Random Projections

We present cancelability results on the clean images from the ND dataset obtained as explained in Section VIII-B3. The iris region obtained after segmentation was unwrapped into a rectangular image of size 10×80. The real parts of the Gabor features were obtained and concatenated to form an iris vector of length 800. We used the random Gaussian matrix in our experiments, though other random matrices mentioned in Section VII-A also gave similar results. In [39], it was shown that separate application of the random projections performed better when compared to the application of a single random projection on the entire iris vector. So we vectorized the real
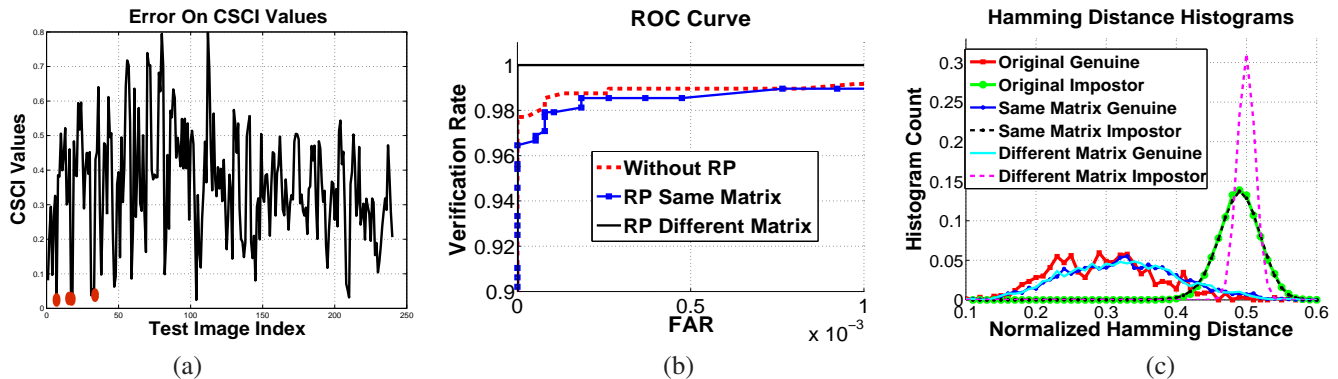
Fig. 12.   (a) Plot of the CSCI values of test images for a random trial on the ND dataset. Red dots indicate the wrongly classified images. Observe that the wrongly classified images have low CSCI values and hence the corresponding vectors are not sparse. (b) ROC characteristics for the ND dataset. The Same Matrix performance is close to the performance without cancelability . Using different matrices for each class gives better performance. (c) Comparison of the distribution of the Genuine and Impostor normalized Hamming distances for the original and transformed patterns.

part of the Gabor features of each sector of the iris image, applied the random projections, and then concatenated the random projected vectors to obtain our cancelable iris biometric. We applied either the same random Gaussian matrix for all the users or different random matrices for different users to obtain the RP "Same Matrix" and "Different Matrix" vectors, respectively. Having obtained the random vectors from the Gabor features of the iris image, we performed the sparsity-based recognition algorithm described in Section III. We present the Receiver Operating Characteristic (ROC) curves and the Hamming distance distributions in the subsections below.

1) Recognition Performance: Fig. 12(b) plots the ROC characteristics for the iris images in the ND dataset for the original and transformed iris patterns. As demonstrated, using different matrices for each class performs better than using the same matrix for all classes. In the "Different Matrix" case, we assumed that the user provided the correct matrix assigned to him. So the performance exceeds even the original performance as class specific random projections increases the interclass distance, still retaining the original intra-class distance. In Fig. 12 (c), we compare the distribution of the genuine and impostor normalized Hamming distance for the original and transformed iris patterns. We can observe that the distribution of the genuine Hamming distance remains almost the same after applying the random projections. The original and Same Matrix cases have similar impostor Hamming distance distributions. However the Different Matrix case has an impostor distribution that is more peaked and farther from the genuine distribution, indicating superior performance.

2) Normalized Hamming distance comparison between the original and the transformed patterns: In this section, we quantify the similarity between the original and the random projected iris vectors. From the original and transformed iris vectors, iris codes are computed by allocating two bits for each Gabor value. The first bit is assigned one if the real part of the Gabor feature is positive and zero otherwise. The second bit is assigned a value of one or zero in a similar manner based on the imaginary part of the Gabor feature. The normalized Hamming distance between the iris codes

is used as the measure of similarity. In Fig. 13(a), we plot the normalized Hamming distance between the iris codes of the original and the transformed iris vectors for the "Same Matrix" and "Different Matrix" cases, respectively. Ideally we want the two iris codes to be independent, hence the normalized Hamming distance should be 0.5. The figure shows that the histogram of the Hamming distance peaks at 0.5, empirically verifying that the random projected iris vectors are significantly different from the originals ones. Hence it is not possible to extract the original iris codes from the transformed version, thereby proving the non-invertibility property of our transformation.

Table V provides the statistics of the normalized Hamming distance between the original and the transformed iris vectors. As can be seen, the mean of the normalized Hamming distance is very close to 0.5 with a very low standard deviation.

TABLE III
STATISTICS OF THE NORMALIZED HAMMING DISTANCE.

Methods	Mean	Standard Deviation
Without RP	0	0
Same Matrix	0.5002	0.0123
Different Matrix	0.4999	0.013
Dictionary Permutations	0.4913	0.0254
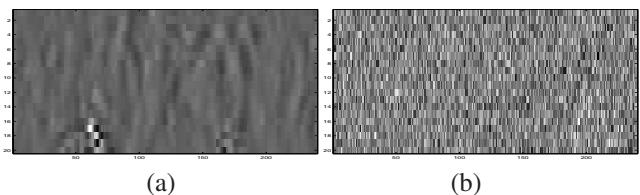
Fig. 14.   (a) Gabor features of the original iris image. (b) Gabor features of the recovered iris image from the cancelable patterns in the dictionary and a randomly generated projection matrix.

3) Non-Invertibility Analysis of Cancelable Templates using Random Projections: In this section, we consider the recovery of original iris patterns from the cancelable templates, using varying levels of information about the dictionary and the projection matrix Φ. We consider two methods, one based on minimizing the squared error and the other based on compressive sensing techniques. As before, we consider eighty

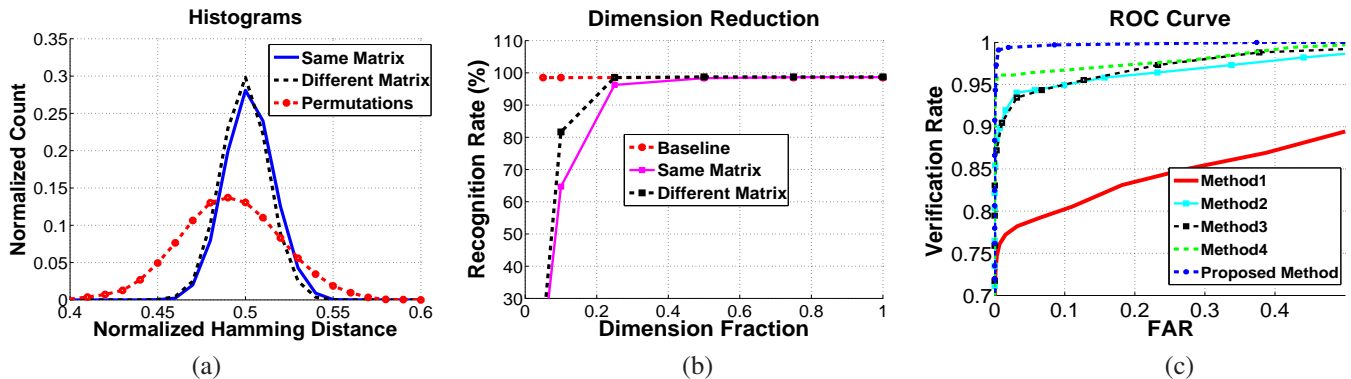Fig. 13. (a) Plot of the histograms of the Normalized Hamming Distance between the original and transformed vectors. Note that the histogram peaks around 0.5 indicating that the original and transformed iris codes are significantly different. (b) Plot of the recognition rate with dimension reductions for the ND dataset. Note that the performance remains the same up to 30% of the original dimension. (c) ROC plots for video based iris recognition. Method 1 treats each frame in the video as a different probe. Method 2 averages all the frames in the probe video. Methods 3 and 4 use the average and minimum of all the pair wise Haming distance between the frames of the probe and gallery videos respectively. The Proposed Method uses CSCI as the matching score. Note that the introduced quality based matching score outperforms the existing fusion schemes, which do not incorporate the quality of the individual frames in the video.

classes from the ND-IRIS-0405 dataset with fifteen images per class. Twelve images per person for the training set and the remaining for the test vectors. We apply the same random projections Φ for each class with a dimension reduction of 40% to form the cancelable patterns. Hence, we have the $\mathbf{a} = \Phi\mathbf{D}\mathbf{y}$, where $\mathbf{a}$ is the cancelable template and $\mathbf{y}$ is the original iris pattern. We consider two methods for reconstructing the original patterns from cancelable patterns. They are explained below.

1) Least Square solution - From equation (23) in the presence of additive noise, the original template can be recovered by minimizing the following squared error.

$$\hat{\mathbf{y}} = \arg\min_{\mathbf{y}} \|\mathbf{a} - \Phi\mathbf{y}\|_2^2$$

2) Compressive Sensing based solution - Since Φ is a random Gaussian matrix having good RIP, one possible way of reconstructing the iris patterns is by solving the following L1 minimization problem.

$$\hat{\mathbf{y}} = \arg\min_{\mathbf{y}} \| \mathbf{y} \|_1 \quad \text{s. t.} \quad \|\mathbf{a} - \Phi\mathbf{y}\|_2 \leq \varepsilon'. \quad (29)$$

We computed the error in reconstruction of the original patterns and the recognition rate on the reconstructed patterns for different levels of information known about the cancelable template dictionary and the random projection matrix Φ. The results are shown in Table V. As can be observed, the recognition performance is close to chance when either the random matrix or the dictionary entries are not known. Even when the random matrix and the dictionary entries are fully known, the recognition performance on the reconstructed template is significantly lower than that on the original iris templates. This result empirically verifies that it is difficult to extract significant information about the original iris templates from the cancelable ones.

In Fig. 14, we display the Gabor features of one of the iris images in the dictionary and the corresponding recovered pattern. As can be observed, the recovered pattern appears as random noise and does not contain any of the information in the original iris pattern.

TABLE IV
RECONSTRUCTION ERROR AND RECOGNITION RATE KNOWING THE EXACT CANCELABLE TEMPLATE AND FRACTION OF ENTRIES IN THE PROJECTION MATRIX

Method	Metric %	Fraction Of Correct Values					
		0	.2	.4	.6	.8	1
LS	Recon. Error	50	49	49	49	49	49
	Recog. Rate	2.9	2.08	2.08	.42	.83	.83
CS	Recon. Error	49	46	42	38	32	22
	Recog. Rate	1.67	2.08	3.33	7.92	24.58	59.17

TABLE V
RECONSTRUCTION ERROR AND RECOGNITION RATE KNOWING THE EXACT PROJECTION MATRIX AND FRACTION OF ENTRIES IN THE CANCELABLE TEMPLATE

Method	Metric (%)	Fraction Of Correct Values					
		0	.2	.4	.6	.8	1
LS	Recon. Error	49	49	49	49	49	49
	Recog. Rate	1.25	2.08	1.25	.83	1.25	2.5
CS	Recon. Error	49	48	46	43	38	22
	Recog. Rate	1.25	1.67	1.25	1.67	9.17	57.50

4) Effect of dimension reduction: In Fig. 13(b), we demonstrate the robustness of random projections to reduction in the original dimension of the feature vector. The random projected vectors retain their original performance for up to 30% reduction in the original dimension for both the same and different matrix cases. Dimension reduction further strengthens the non-invertibility of our transformation as there will be infinite possible iris vectors corresponding the reduced dimension random vectors obtained by RP.

5) Comparison with Salting: In Table. VI, we present the recognition rates and the corresponding mean Hamming distance for the salting method proposed in [29] for various noise levels. The best recognition rate and the best Hamming distance for the Salting method are 96.6% and 0.494 respectively. For RP Same Matrix case, we obtained a recognition rate of 97% at a Hamming distance of .497. Thus both the recognition performance and security (non-invertibility) are higher for RP when compared to the Salting method.

TABLE VI
COMPARISON WITH SALTING METHOD. THE RECOGNITION RATE(RR)
AND MEAN HAMMING DISTANCE (HD) ARE PROVIDED FOR THE SALTING
AND SRP METHODS. THE RECOGNITION RATE OBTAINED USING SRP IS
HIGHER THAN THAT OF THE SALTING METHOD. ALSO SRP GIVES MEAN
HAMMING DISTANCE CLOSER TO .5 WHEN COMPARED TO THE SALTING
METHOD.

Quantity	Salting			Same	Different	Permutations
RR(%)	94.2	96.6	94.0	97	100	100
HD	0	.491	.494	.497	.50	.483

D. Cancelability Results using Random Permutations

To evaluate the performance of the proposed cancelable method using dictionary permutations, we consider the three possible scenarios on the clean images from the ND dataset. In the first case, the user provides the iris image and the correct hash code. In this case, the recognition performance was the same as that of the original method on the ND dataset, which is 99.17%. In the second case, the user provides the iris image but a wrong hash code. Here the recognition performance dropped to 2%, which is only slightly better than chance. This is equivalent to the case when a hacker illegally obtains the iris image of a valid user and tries to gain access into the system with a guess about the hash code. The low recognition performance clearly reflects the additional security introduced by the permutations, as a hacker needs to now have not only the iris image but also the hash code of a valid user to gain access. In the third experiment, we found the closeness between the Gabor features of the original iris images and the new feature vectors obtained by permutations of the Gabor features in the dictionary. As before, the normalized Hamming distance between the iris codes obtained from these vectors is used as the measure of similarity. We plot the histogram of the normalized Hamming distance between the original and the randomly permuted iris vectors in Fig. 13(a). The mean and standard deviation of the Hamming distance histogram are indicated in the last row of the Table. V. Note that the mean is close to .5, indicating that the permutations differ significantly different from the original iris images. Even if a hacker can use the dictionary from the application database, he or she will be unable to extract information about the original iris images without knowing the hash code of each user.

E. Results on Iris Videos

In this section, we present the results on the MBGC videos [45]. Given the thirty classes, we used twenty eight classes that contained atleast five good images in our experiments. We hand picked five clean images from the iris videos in the training set which formed the dictionary. In the test videos, batches of five frames were given as a probe to our algorithm. Using twenty eight available videos and sixty frames from each test video, we could form three hundred and thirty six probes. We did only a basic segmentation of the iris and pupil using the Masek's code, as before. Also, we did not remove the poorly segmented iris images manually before performing the recognition algorithm.

We compare the performance of our algorithm with four other methods. The ROC plots for the different methods are displayed in Fig. 13(c). In Method 1, we consider each frame of the video as a different probe. It gave the worst performance, indicating that using multiple frames available in a video can improve the performance. Method 2 averages the intensity of the different iris images. Though it performs well when the images are clean, a single image which is poorly segmented or blurred could affect the entire average. In Methods 3 and 4, all possible pair wise Hamming distances between the video frames of the probe videos and the gallery videos belonging to the same class are computed. Method 3 uses the average of these Hamming distance as the score. In Method 4, the minimum of the pairwise Hamming distance was used as the score. In the proposed method, the CSCI values were computed for each class for each probe video and the probe video is assigned to the class having the highest CSCI value. For a fair comparison of the proposed quality measure in videos, we did not reject any of the frames. Observe that our method performs better than other methods. One of the reasons for the superior performance could be the fact that we are incorporating the quality of the different frames while computing the CSCI. Frames which are poorly segmented or blurred will have a low SCI value and hence will not affect the score significantly. In all the other methods, the image quality was not effectively incorporated into the matching score, so all frames are treated equally irrespective of their quality.

F. Effect of Non-Linear Deformations

Thornton *et. al.* [49] have studied the in-plane nonlinear deformations possible in iris images. They found that the movement of the pupil is only approximately linear, which leads to minor relative deformations to the normalized iris patterns.

For BPDN in Section III.A to identify the correct class, the fitting error of the true class should be much lower than that due to the other classes. The non-linear deformations could increase the fitting error for the true class and affect the performance of our algorithm. To analyze this issue empirically, we represent each sector of the test feature as a linear combination of the training images of each class separately. The median of the fitting error for the true class and the other classes is plotted in Fig. 15 for the different sectors. As can be observed, the fitting error for the true class is significantly lower than that using the other classes. Also the number of training images in the true class is far fewer than the total number of images in the gallery, thereby making the coefficient vector sparse. Empirically, we can thus verify that the proposed method should be returning the correct class even with small non-linear deformations.

Also, the non-linear deformations will differ for different sectors of the iris. If some sectors have lower deformation and others have higher deformation, the SCI values will be greater for the lesser deformed sectors and lower for the rest. The sectors with lower deformation will be correctly recognized by the Sparse Representation based method. The introduced Bayesian Fusion framework will give greater weight to the evidence from the lesser deformed sectors, thereby obtaining
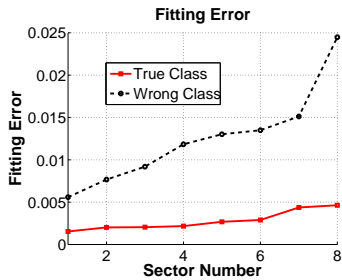
Fig. 15. Plot of the fitting error for the true class and the wrong classes. Note that the true class gives a much lower fitting error on the test iris images. Also, the number of images in the true class is far fewer than the total number of images in the gallery. So the sparsest solution, which gives the least fitting error corresponds to non zero coefficients only for the true class, in most cases.

the correct class label. So the performance of our method will be affected only when all sectors of the iris have severe non-linear deformations. We are currently investigating how to handle such severe non-linear deformations within our framework.

G. Computational Complexity

In this section, we analyze the complexity of our algorithm and compare it with the Masek's algorithm for iris recognition. Let us assume that we have L classes, with n training images per class. Let each feature vector be of length N. Let us assume that the number of sectors is n_s. In our experiments, it is fixed as 8. So the length of each sector of the feature vector is $\frac{N}{n_s}$. The size of the dictionary in each sparse representation is $\frac{N}{n_s} \times nL$.

Numerous methods have been proposed for solving the ℓ_1 minimization problem [50]. For instance, the complexity of solving an ℓ_1 minimization problem using a Homotopy based method is an $\mathcal{O}(\frac{N}{n_s}^2 + \frac{N}{n_s}nL)$ operation. So the total computational complexity of finding the sparse representation of all the sectors is $\mathcal{O}(n_s(\frac{N}{n_s}^2 + \frac{N}{n_s}nL))$. Let $\hat{A}$ be the number of different alignments used. So the total complexity of our algorithm is $\mathcal{O}(n_s\hat{A}(\frac{N}{n_s}^2 + \frac{N}{n_s}nL))$. Since $\hat{A}$ and s_l are constants in our algorithm, the total complexity of the proposed method is $\mathcal{O}(N^2 + NnL)$. Note that computing CSCI given the SCI of the different sectors is an $\mathcal{O}(L\hat{A})$ operation. So the complexity of our algorithm is linear in the number of classes and quadratic in the length of the feature vector. The computational complexity of the Masek's algorithm is $\mathcal{O}(NnL)$, which is linear in the number of classes and the dimension of the feature vector. So our algorithm is computationally more expensive than the Masek's algorithm and the difference in complexity increases with the feature dimension. However, since the most computationally expensive component in our algorithm is the ℓ_1 minimization step, it is possible to make the proposed algorithm faster using greedy pursuit algorithms such as CoSaMP [51].

IX. CONCLUSION

In this paper, we proposed a unified approach for iris image selection and recognition, which has numerous ad-

vantages over existing techniques when sufficient number of training samples are available. Our experiments indicate that the selection algorithm can handle common distortions in iris image acquisition like blur, occlusions and segmentation errors. Unlike the existing feature based and score based fusion algorithms for iris recognition from video, we introduced a quality-based matching score and demonstrated its superior performance on the MBGC iris video dataset. Furthermore, we incorporated random projections and random permutations into the proposed method to prevent the compromise of sensitive biometric information of the users. Currently, we are exploring ways to improve the quality of the poorly acquired images, rather than rejecting them using compressed sensing techniques [33], [41]. We are also investigating ways to develop a minimal set of training images spanning the testing image space, which can reduce the memory and computation requirements of our algorithm. The possibility of adapting algorithms such as elastic net [52] is also being explored.

REFERENCES

[1] K. W. Bowyer, K. Hollingsworth, and P. J. Flynn, "Image understanding for iris biometrics: A survey," *Computer Vision and Image Understanding*, vol. 110, no. 2, pp. 281–307, 2008.

[2] J. Daugman, "Probing the Uniqueness and Randomness of Iris Codes:Results from 200 Billion Iris Pair Comparisons," *Proceedings of the IEEE*, vol. 94, no. 11, 2006.

[3] E. M. Newton, P. J. Phillips, "Meta-Analysis of Third-Party Evaluations of Iris Recognition," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 39, no. 1, pp. 4–11, 2009.

[4] J. Wright, A. Y. Yang, A. Ganesh, S. S. Sastry, and Y. Ma, "Robust Face Recognition via Sparse Representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 2, pp. 210–227, 2009.

[5] H. Proenca and L. A. Alexandre, "Iris segmentation methodology for non-cooperative recognition," In *IEEE Proceedings on Vision, Image and Signal Processing*, vol. 153, pages 199-205, 2006.

[6] A. K. Jain, K. Nandakumar, and A. Nagar, "Biometric Template Security," *EURASIP Journal on Advances in Signal Processing, Special issue on Biometrics*, vol. 2008, no. 113, pp. 1–17, 2008.

[7] R.M. Bolle, J.H. Connel and N.K. Ratha, "Biometrics Perils and Patches," *Pattern Recognition*, vol. 35, no. 12, pp. 2727–2738, 2002.

[8] N.K. Ratha, J.H. Connel and R.M. Bolle "Enhancing security and privacy in biometrics-based authentication systems," *IBM Systems Journal*, vol. 40, no. 3, pp. 614–634, 2001.

[9] J. Daugman, "High confidence visual recognition of persons by a test of statistical independence," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15, pp. 1148-1161, 1993.

[10] R. P. Wildes, "Iris recognition: An emerging biometric technology," *Proceedings of the IEEE* vol. 85, no. 9, pp. 1348–1363, 1997.

[11] Y. Chen, S. C. Dass and A. K. Jain, "Localized iris image quality using 2-D wavelets," *Springer Lecture Notes in Computer Science 3832: International Conference on Biometrics*, pages 373-381, 2006.

[12] N. D. Kalka, J. Zuo, N. A. Schmid and B. Cukic, "Image quality assessment for iris biometric," In *SPIE 6202: Biometric Technology for Human Identification*, pages 6202, 2006.

[13] H. Proenca and L. A. Alexandre, "A method for the identification of noisy regions in normalized iris images," *International Conference on Pattern Recognition*, pp. 405–408, 2006.

[14] X.-D. Zhu, Y.-N. Liu, X. Ming, and Q.-l. Cui, "Quality Evaluation Method of Iris Images Sequence Based on Wavelet Coefficients in Region of Interest," *The Fourth International Conference on Computer and Information Technology (CIT)*, pp. 24-27, 2004.

[15] L. Ma, T. Tan, Y. Wang, D. Zhang, "Personal Identification Based on Iris Texture Analysis," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25 , no. 12, pp. 1519–1533, 2003.

[16] S. Rakshit, D. Monro, "Iris image selection and localization based on analysis of specular reflection," *IEEE Workshop on Signal Processing Applications for Public Security and Forensics, SAFE*, pp. 1–4, 2007.

[17] K. P. Hollingsworth, K. W. Bowyer and P. J. Flynn, "Image Averaging for Improved Iris Recognition," *Lecture Notes in Computer Science*, vol. 5558, 2009.

[18] Y. Du, "Using 2D log-Gabor spatial filters for iris recognition," *SPIE Biometric Technology for Human Identification III*, vol. 62020, 2006.

[19] L. Ma, T. Tan, Y. Wang, D. Zhang, "Efficient iris recognition by characterizing key local variations," *IEEE Transactions on Image Processing*, vol. 13, no. 6, pp. 739750, 2004.

[20] E. Krichen, L. Allano, S. Garcia-Salicetti, B. Dorizzi, "Specific texture analysis for iris recognition," Lecture Notes in Computer Science, vol. 3546, pp. 2330. Springer, Heidelberg, 2005.

[21] N. A. Schmid, M. V. Ketkar, H. Singh, B. Cukic, "Performance analysis of iris based identification system at the matching score level," *IEEE Transactions on Information Forensics and Security* vol. 1, no. 2, pp. 154-168, 2006.

[22] C. Liu, M. Xie, "Iris recognition based on DLDA," *International Conference on Pattern Recognition*, pp. 489-492, 2006.

[23] K. Roy, P. Bhattacharya, "Iris recognition with support vector machines," *International Conference on Biometrics*, pp. 486492, 2006.

[24] G. I. Davida, Y. Frankel and B. J. Matt, "On enabling secure applications through off-line biometric identification," *IEEE Symposiam on Security and Privacy*, pp. 148-157, 1998.

[25] F. Hao, R. Anderson, and J. Daugman, "Combining crypto with biometrics effectively," *IEEE Transactions on Computers*, vol. 55, no. 9, pp. 1081–1088, 2006.

[26] S. Kanade, D. Petrovska-Delacretaz and B. Dorizzi, "Cancelable iris biometrics and using error correcting codes to reduce variability in biometric data," *Computer Vision and Pattern Recognition (CVPR)*, 2009.

[27] A. Juels and M. Wattenberg, "A fuzzy commitment scheme," *ACM Conference on Computers and Communications Security*, pp. 28–36, 1999.

[28] A.B.J. Teoh, A. Goh and D.C.L. Ngo, "Random Multispace Quantization as an Analytic Mechanism for BioHashing of Biometric and Random Identity Inputs," *IEEE Transactions on Pattern Analysis and Machine Intelligence*,vol. 28, no. 12, pp. 1892–1901, 2006.

[29] J. Zuo, N.K. Ratha and J.H. Connell, "Cancelable Iris Biometric," *Proc. International Conference on Pattern Recognition*, pp. 1–4, 2008.

[30] D. L. Donoho and X. Huo, "Uncertainty principle and ideal atomic decomposition," *IEEE Transactions on Information Theory*, vol. 47, no. 7, pp. 2845–2862, 2001.

[31] D. L. Donoho and M. Elad, "On the Stability of the Basis Pursuit in the Presence of Noise," *EURASIP Signal Processing Journal*, vol. 86, no. 3, pp. 511-532, 2006.

[32] E. Candès, J. Romberg and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications in Pure and Applied Mathematics*, vol. 59, pp. 1207-1223, 2006.

[33] R. Baraniuk, "Compressive sensing," *IEEE Signal Processing Magazine*, vol. 24, No. 4, pp. 118-121, 2007.

[34] S. Chen, D. Donoho, and M. Saunders, "Atomic Decomposition by Basis Pursuit," *SIAM J. Sci. Comp.*,vol. 20, no. 1, pp. 33-61, 1998.

[35] J. A. Tropp, A. C. Gilbert, "Signal Recovery From Random Measurements Via Orthogonal Matching Pursuit," *IEEE Transactions on Information Theory*, vol. 53, no. 12, pp. 4655-4666, 2007.

[36] P. J. Phillips, "Matching pursuit filters applied to face identification," *IEEE Transactions on Image Processing*, vol. 7, no. 8, pp. 150-1164, 1998.

[37] J. K. Pillai, V. M. Patel, R. Chellappa, "Sparsity Inspired Selection and Recognition of Iris Images," *Third IEEE International Conference on Biometrics - Technology And Systems (BTAS)*, 2009.

[38] A.B.J. Teoh, and C.T. Yuang, "Cancelable biometrics realization with multispace random projections," *IEEE Transactions on Systems, Man, and Cybernetics-Part B: Cybernetics* , vol. 37, no. 5, pp. 1096–1106, 2007.

[39] 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*, 2010.

[40] H. Rauhut, K. Schnass, and P. Vandergheynst, "Compressed Sensing and Redundant Dictionaries," *IEEE Transactions on Information Theory*, vol. 54, no. 5, pp. 2210–2219, 2008.

[41] E. Candes, J. Romberg and T. Tao, "Robust Uncertainty Principles: Exact Signal Reconstruction from Highly incomplete Frequency Information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, 2006.

[42] D. L. Donoho, "High-Dimensional centrally symmetric polytopes with neighborliness proportional to dimension," *Discrete and Computational Geometry*, vol. 35, no. 4, pp. 617–652, 2006.

[43] J. D. Blanchard, C. Cartis and J. Tanner, "The restricted isometry property and ℓ_q-regularization: Phase transition for sparse approximation," submitted.

[44] K. W. Bowyer and P. J. Flynn,"The ND-IRIS-0405 Iris Image Dataset," Notre Dame CVRL Technical Report.

[45] P. J. Phillips, P. J. Flynn, J. R. Beveridge, W. T. Scruggs, A. J. O'Toole, D. S. Bolme, K. W. Bowyer, B. A. Draper, G. H. Givens, Y. M. Lui, H. Sahibzada, J. A. Scallan, S. Weimer, "Overview of the Multiple Biometrics Grand Challenge," The Third IAPR/IEEE Conference on Biometrics, 2009.

[46] E. van den Berg and M. P. Friedlander, "Probing the Pareto frontier for basis pursuit solutions," *SIAM J. Sci. Comp*, vol. 31, no. 2, pp. 890-912, 2008.

[47] L. Masek, P. Kovesi, "MATLAB Source Code for a Biometric Identification System Based on Iris Patterns," The University of Western Australia, 2003.

[48] P. J. Phillips, "FRGC and ICE Workshop," NRECA Conference Facility, 2006.

[49] J. Thornton, M. Savvides, and B. V. K. V. Kumar, "A Bayesian Approach to Deformed Pattern Matching of Iris Images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 596–606, 2007.

[50] A. Yang, A. Ganesh, S. Sastry, Y. Ma, "Fast L1-Minimization Algorithms and An Application in Robust Face Recognition : A Review," Technical Report No. UCB/EECS-2010-13, 2010.

[51] D. Needell and J. A. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Appl. Comp. Harmonic Anal.*, vol. 26, pp. 301–321, 2008.

[52] H. Zou and T. Hastie, "Regularization and Variable Selection via the Elastic Net," *Journal of the Royal Statistical Society*, vol. 67, no. 2, pp. 301-320, 2005.

PLACE PHOTO HERE

Jaishanker K Pillai received his B. Tech. (Distinction) degree in Electronics and Communication Engineering from National Institute Of Technology, Calicut and his M.E. (Distinction) degree in System Science and Automation from Indian Institute of Science, India in 2006 and 2008, respectively. He is currently pursuing the Ph.D. degree in electrical and computer engineering at the University of Maryland, College Park. His research interests include computer vision, pattern recognition, signal processing and biometrics. Mr. Pillai was awarded the gold medal for the best outgoing student for his Bachelors and Masters degrees respectively. He is also the recipient of the University of Maryland James Clarke Graduate Fellowship for 2008-10. He is a student member of IEEE.

PLACE PHOTO HERE

Vishal M. Patel Vishal M. Patel (M'01) received the B.S. degrees in electrical engineering and applied mathematics (with honors), the M.S. degree in applied mathematics from North Carolina State University, Raleigh, NC, in 2004 and 2005, respectively, and the Ph.D. degree in electrical engineering from the University of Maryland, College Park, MD, in 2010.

His research interests include compressed sensing, biometrics, radar imaging and pattern recognition.

He is a member of Eta Kappa Nu, Pi Mu Epsilon, Phi Beta Kappa and a member of SIAM.

PLACE PHOTO HERE

Rama Chellappa received the B.E. (Hons.) degree from University of Madras, India, in 1975 and the M.E. (Distinction) degree from Indian Institute of Science, Bangalore, in 1977. He received M.S.E.E. and Ph.D. Degrees in Electrical Engineering from Purdue University, West Lafayette, IN, in 1978 and 1981 respectively. Since 1991, he has been a Professor of Electrical Engineering and an affiliate Professor of Computer Science at University of Maryland, College Park. He is also affiliated with the Center for Automation Research (Director) and the Institute for Advanced Computer Studies (Permanent Member). In 2005, he was named a Minta Martin Professor of Engineering. Prior to joining the University of Maryland, he was an Assistant (1981-1986) and Associate Professor (1986-1991) and Director of the Signal and Image Processing Institute (1988-1990) at University of Southern California, Los Angeles. Over the last 29 years, he has published numerous book chapters, peer-reviewed journal and conference papers. He has co-authored and edited books on MRFs, face and gait recognition and collected works on image processing and analysis. His current research interests are face and gait analysis, markerless motion capture, 3D modeling from video, image and video-based recognition and exploitation, compressive sensing, and hyper spectral processing. Prof. Chellappa served as the associate editor of four IEEE Transactions, as a Co-Editor-in-Chief of Graphical Models and Image Processing and as the Editor-in-Chief of IEEE Transactions on Pattern Analysis and Machine Intelligence. He served as a member of the IEEE Signal Processing Society Board of Governors and as its Vice President of Awards and Membership. He is serving a two-year term as the President of IEEE Biometrics Council. He has received several awards, including an NSF Presidential Young Investigator Award, four IBM Faculty Development Awards, an Excellence in Teaching Award from the School of Engineering at USC, two paper awards from the International Association of Pattern Recognition. He received the Society, Technical Achievement and Meritorious Service Awards from the IEEE Signal Processing Society. He also received the Technical Achievement and Meritorious Service Awards from the IEEE Computer Society. At University of Maryland, he was elected as a Distinguished Faculty Research Fellow, as a Distinguished Scholar-Teacher, received the Outstanding Faculty Research Award from the College of Engineering, an Outstanding Innovator Award from the Office of Technology Commercialization and an Outstanding GEMSTONE Mentor Award. He is a Fellow of the IEEE, the International Association for Pattern Recognition and Optical Society of America. He has served as a General the Technical Program Chair for several IEEE international and national conferences and workshops. He is a Golden Core Member of the IEEE Computer Society and served a two-year term as a Distinguished Lecturer of the IEEE Signal Processing Society.

PLACE PHOTO HERE

Nalini K. Ratha received his Ph. D. from the Department of Computer Science at Michigan State University and B.Tech in Electrical Engineering and M. Tech in Computer science and Engineering from Indian Institute of Technology, Kanpur. Currently, he is a Research Staff Member at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York where he leads the biometrics research efforts in building efficient biometrics systems. In addition to more than 80 publication in peer reviewed journals and conferences, and co-inventor on 12 patents, he has co-edited two books entitled Advances in Biometrics: Sensors, Algorithms and Systems and Automatic Fingerprint Recognition Systems published by Springer and co-authored a popular text book entitled A guide to Biometrics published by Springer in 2003. In the past, he has been associated with several leading biometrics conference including general co-chair of IEEE AutoID 02 and SPIE Conf. on Biometrics in Human Identification 2004 2005. More recently, he was co-chair of an associated theme on biometrics at ICPR 2006 and ICPR 2008 and also the founding co-chair of CVPR workshop on Biometrics 2006-2010, and founding co-program chair of IEEE Biometrics Theory, Applications and Sytems (BTAS 07 – BTAS 09). He served on the editorial board of Pattern Recognition journal and IEEE Trans. on Image Processing. He was a guest co-editor of a special issue on biometrics for IEEE Trans. on SMC-B and also for IEEE Trans. on Information Forensics and Security. Currently, he serves on the editorial board of IEEE Trans. on Pattern Analysis and Machine Intelligence and IEEE Trans. on Systems, Man and Cybernetics- Part B. He has received several patent awards and a Research Division award and two Outstanding Technical Innovation Award at IBM. He continues to teach at Cooper Union and NYU-Poly as an Adjunct professor for last several years. He is a Fellow of IEEE, Fellow of IAPR and Senior Member of ACM. His current research interests include biometrics, computer vision, pattern recognition and special purpose architecture for computer vision systems.