

# Graphical Model Approach to Iris Matching Under Deformation and Occlusion

R. Kerekes B. Narayanaswamy J. Thornton M. Savvides B. V. K. Vijaya Kumar  
ECE Department, Carnegie Mellon University  
Pittsburgh, PA 15213

## Abstract

*Template matching of iris images for biometric recognition typically suffers from both local deformations between the template and query images and large occlusions from the eyelid. In this work, we model deformation and occlusion as a set of hidden variables for each iris comparison. We use a field of directional vectors to represent deformation and a field of binary variables to represent occlusion. We impose a probability distribution on these fields using a lattice-type undirected graphical model, in which the graph edges represent interdependencies between neighboring iris regions. Gabor wavelet-based similarity scores and intensity statistics are used as observations in the model. Loopy belief propagation is applied to estimate the conditional distributions on the hidden variables, which are in turn used to compute final match scores. We present underlying theory as well as experimental results from both the CASIA iris database and the database provided for the Iris Challenge Evaluation (ICE). We show that our proposed method significantly improves recognition accuracy on these datasets over existing methods.*

## 1. Introduction

The iris is the colored region of the eye bounded by the pupil and white sclera. The intensity pattern in the iris is believed to be detailed enough to reliably discriminate between individuals and remains stable over the human lifetime [12]. Iris recognition has recently emerged as an important component of the field of biometric recognition. The work of Daugman [3] and others has demonstrated that the human iris pattern, much like a fingerprint, is a useful biometric for determining and/or verifying the identity of a person.

In most iris recognition algorithms, an image of the iris is first segmented by finding two approximately concentric circles that define the inner and outer boundaries of the pattern region. Once this region is determined, it is subsequently “unwrapped” from Cartesian to polar coordinates. Such an unwrapping is illustrated in Fig. 1. The resulting

image is used as the basis for feature analysis and comparison. Because of local muscle contractions and other biological phenomena, two iris patterns captured from the same eye are often spatially deformed relative to one another. In addition, there almost always exist significant portions of the iris pattern which are occluded by the eyelid as shown in Fig. 2, and the amount of occlusion can vary between samples. A robust iris recognition algorithm must be able to handle these non-linear deformations and local occlusions.

Recent work by Thornton *et al.* [11] has shown that iris recognition performance can be significantly improved by partitioning the template iris pattern into multiple regions and individually matching these smaller regions to the query image. The approach in [11] iteratively estimates each local displacement vector by holding the displacements of all other regions fixed and applying the Iterated Conditional Modes (ICM) algorithm [6]. The resulting deformation estimate is then used to assign an overall match score. The authors show that imposing such a probabilistic model results in better performance than simply estimating the deformation without any constraints. Probabilistic methods for estimating deformation fields have been proposed in [15] and [9], and similar methods for estimating occlusion have been reported in [14].

We propose an approach that includes the following modifications: (1) estimating occluded regions jointly and probabilistically with the deformations rather than separately and deterministically, and (2) using soft information from the estimated distribution, rather than a single point estimate, to compute a Bayesian score. Such an approach avoids discarding most of the information returned by the region matching algorithm and thus has the potential to yield an improved comparison score. We show that our proposed method results in a large reduction in error rates over a baseline algorithm on two separate iris databases.

## 2. Background

### 2.1. Database description

We used two databases of iris images to test our method. The first one is the CASIA (Chinese Academy of Sciences

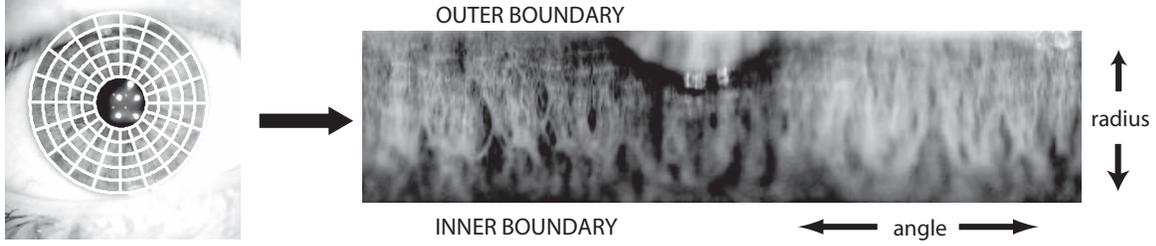


Figure 1. Example of fitting a polar coordinate system to an iris pattern for subsequent unwrapping.

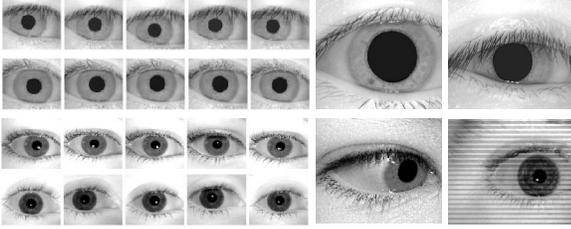


Figure 2. Sample iris images from the CASIA database (*top*) and ICE database (*bottom*). The larger images are representative of difficult images within each database.

Institute of Automation) iris database ver. 1.0 [1], which consists of 108 classes with 7 samples from each class for a total of 756 images. Each image is  $280 \times 320$  pixels in size. The second and larger database is the one supplied for the NIST Iris Challenge Evaluation (ICE) [2] (Phase I, Experiments 1 and 2). The ICE databases consist of 124 classes (Experiment 1) and 120 classes (Experiment 2) with between 1 and 31 images of size  $480 \times 640$  in each class. Experiment 1 contains a total of 1,425 images, while Experiment 2 contains a total of 1,528 images. Sample images from each database are shown in Fig. 2. The reasons for choosing these two databases are as follows:

1. Some parameters in our model need to be learned from training data. These are learned on the CASIA database and then used on both CASIA and ICE, showing that these parameters are not database specific.
2. The two databases have, among other properties, different image resolutions and imaging artifacts. Our algorithm is shown to work well on both, demonstrating its general applicability.

## 2.2. Data preprocessing

The iris pattern in each image is first segmented by detecting two circular boundaries, a larger circle separating the outer iris from the sclera, and a smaller circle separating the inner iris from the pupil. We use the segmentation

algorithm described in [11]. Once these boundaries are determined, a polar coordinate system is fit to the iris region, as shown in Fig. 1. The image resulting from unwrapping this polar coordinate system into an  $M \times N$  grid is referred to as the iris plane. The dimensions  $M$  and  $N$  are fixed for a given set of comparisons. Let  $f(\mathbf{x})$  denote the iris plane, where  $\mathbf{x} = [x \ y]^T \in \{1 \dots M\} \times \{1 \dots N\}$  denotes the pixel location.

Iris code images  $\mathbf{c}(\mathbf{x})$  are created from each iris plane by extracting a set of features at every point. We use the feature extraction method described in [3], wherein the iris pattern is projected onto a set of complex-valued 2D Gabor wavelets  $\{g_i(\mathbf{x})\}_{i=1}^N$ . Specifically, each  $g_i(\mathbf{x})$  is of the form

$$g_i(\mathbf{x}) = g_0(\mathbf{R}_{\theta_i} \mathbf{x}; \sigma_{x_i}, \sigma_{y_i}, \Omega_i) \quad (1)$$

where the unrotated wavelet function  $g_0$  is defined as

$$g_0(\mathbf{x}; \sigma_x, \sigma_y, \Omega) = \exp \left[ -\frac{1}{2} \left( \frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} \right) + j2\pi\Omega x \right] \quad (2)$$

In (1),  $\sigma_{x_i}$ ,  $\sigma_{y_i}$ ,  $\theta_i$ , and  $\Omega_i$  are the wavelet parameters of the  $i$ -th wavelet, and  $\mathbf{R}_{\theta}$  is a 2-D rotation matrix of angle  $\theta$ . An array of binary-valued (1 and -1) feature vectors  $\mathbf{c}(\mathbf{x})$  is then computed as

$$\mathbf{c}(\mathbf{x}) = \begin{bmatrix} \text{sgn} \{ \text{Re} [f * g_1(\mathbf{x})] \} \\ \text{sgn} \{ \text{Im} [f * g_1(\mathbf{x})] \} \\ \vdots \\ \text{sgn} \{ \text{Re} [f * g_N(\mathbf{x})] \} \\ \text{sgn} \{ \text{Im} [f * g_N(\mathbf{x})] \} \end{bmatrix} \quad (3)$$

where the  $*$  operator denotes a 2-D discrete convolution. The feature array  $\mathbf{c}(\mathbf{x})$  may be downsampled to a lower spatial resolution (i.e.,  $\mathbf{c}'(\mathbf{x}) = \mathbf{c}(L\mathbf{x})$ ) in order to decrease the number of features and reduce redundancy.

The iris code generation is a shift-invariant process; i.e., if the input iris plane undergoes a certain shift (in polar coordinates), the resulting iris code will be shifted in exactly the same way. The particular Gabor wavelets used (i.e., the scale and rotation parameters plus the values of  $\sigma_x$  and  $\sigma_y$ )

and the resolution (downsampling factor  $L$ ) of the resulting iris code are chosen to maximize the accuracy of the baseline method proposed in [3] (described in Section 4).

### 2.3. Match score array generation

A match score between any template and query image pair is computed as a function of the total Hamming distance between the resulting iris codes. In both the baseline and proposed matching algorithms, we allow for the possibility that the images may be shifted relative to one another in the polar domain; therefore, we compute match scores between the two images at all plausible shifts. Let  $\mathbf{c}_t(\mathbf{x})$  and  $\mathbf{c}_q(\mathbf{x})$  denote the unshifted template and query iris codes, respectively. The match score array  $m(\mathbf{x})$  at each shift vector  $\mathbf{x}$  is then given by the following equation:

$$m(\mathbf{x}) = \frac{1}{|S_t|} \sum_{\mathbf{y} \in S_t} \mathbf{c}_t^T(\mathbf{y}) \mathbf{c}_q(\mathbf{y} - \mathbf{x}), \quad (4)$$

where  $S_t$  is the support of the template iris code and  $|S_t|$  is the size of the template. It can be shown that each computed score in the array in (4) is a function of the Hamming distance between the two codes at that particular shift.

## 3. Iris comparison approach

### 3.1. Deformation model

As a result of pupil dilation, the iris image (and thus the iris code) can undergo a nonlinear deformation between two images from the same class. We approximate this nonlinear deformation as a coarse vector field. Specifically, we first divide the template iris code image into  $N_s$  regions, as shown in Fig. 3. We then assume that when a query image from the same class is compared to the template, each template region will be closely matched to a nearby region in the query image, but that the relative displacement between a pair of corresponding regions may be different from that of the other pairs in the images. For our experiments, we found  $N_s = 36$  to work well.

In order to evaluate the possible alignments, we compute similarity measures between each region in the template and the query image across all plausible shifts. We use the technique of Section 2.3 to generate  $N_s$  separate match score arrays  $\{m_i(\mathbf{x})\}_{i=1}^{N_s}$ , where each  $m_i(\mathbf{x})$  is computed per Eq. (4) but with  $\mathbf{c}_t(\mathbf{x})$  replaced by  $\mathbf{c}_{it}(\mathbf{x})$ , the  $i$ -th region of the template iris code. Only match scores corresponding to a shift of 10 pixels or less in either direction (vertical or horizontal) are retained. Thus for an iris plane partitioned into  $N_s$  regions, this matching scheme generates  $N_s$  different score arrays of size  $21 \times 21$ .

### 3.2. Occlusion metric

For each pixel of the iris plane  $\mathbf{x}$ , we compute an occlusion metric  $\pi(\mathbf{x})$  that measures the likelihood that the pixel belongs to eyelid rather than to the iris pattern. This metric is computed from four local statistics: (1) the mean intensity value in a small neighborhood around the pixel, (2) the standard deviation of the intensity values in the same neighborhood, (3) the percentage of pixels whose intensity is greater than one standard deviation above the mean of the entire iris plane, and (4) the shortest Euclidean distance to the centers of the upper and lower eyelids. A Fisher linear discriminant [4] is then used to generate a single scalar quantity  $\pi(\mathbf{x})$  from this set of four statistics at each pixel. Finally, an overall occlusion metric  $\pi_i$  is computed for each region  $i$  as the mean of all  $\pi(\mathbf{x})$  in that region. The Fisher linear discriminant was trained using 30 iris plane images whose eyelid regions were labeled manually. We used a neighborhood size of  $7 \times 7$  pixels.

### 3.3. Graphical model for deformations and occlusions

If a certain segment in the image has a particular deformation or occlusion state, neighboring segments are more likely to have similar states rather than different ones. Also, images that require large deformations to obtain a good match are less likely to belong to the same class than images that obtain high scores with small deformations. In this section we investigate a method to use such information in the matching process.

We use an undirected lattice-type graphical model, depicted in Fig. 4, to model the interactions between the deformations and occlusions of neighboring regions in the partitioned iris code. Each node  $\mathbf{d}_i$  in the model,  $i = 1, \dots, N_s$ , represents a 2-D discrete-valued shift vector, the true value of which is hidden. The components of  $\mathbf{d}_i$  are the vertical and horizontal shifts (in pixels) of the template region relative to the corresponding query region. The nodes  $\omega_i$ ,

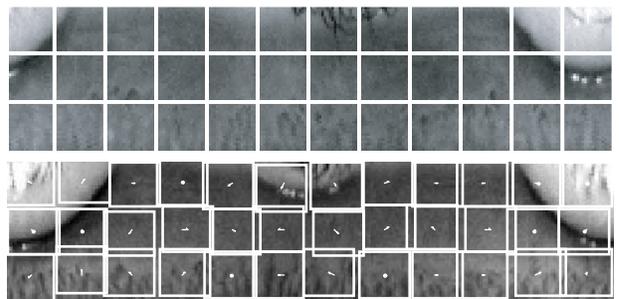


Figure 3. Partitioning of the template iris plane into  $N_s = 36$  subregions (top), and an example of local shifts between template and query images (bottom).

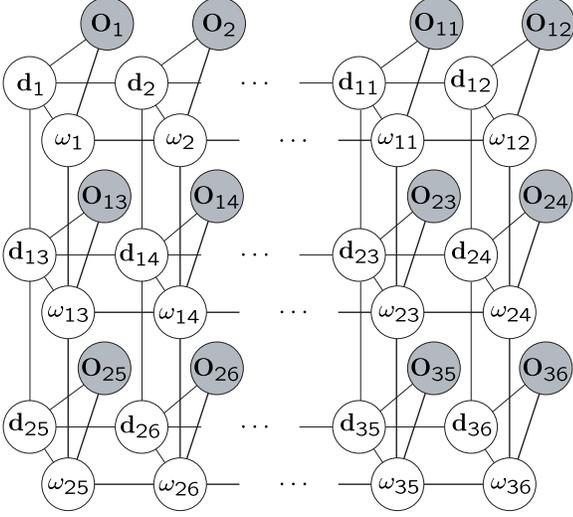


Figure 4. Undirected graphical model used to model the deformation distribution ( $3 \times 12$  configuration shown).

$i = 1, \dots, N_s$ , are hidden binary-valued occlusion variables, where  $\omega_i = 0$  and  $\omega_i = 1$  denote that the region is occluded and unoccluded by the eyelid, respectively. Nodes  $\mathbf{O}_i$  represent the observations, which include the match score arrays  $m_i(\mathbf{x})$  and the occlusion statistic  $\pi_i$  computed for each template region.

In specifying the model, we need to specify potential functions between each pair of connected nodes in the graphical model. The output of any potential function for a given input should correspond to the “goodness” of the input, i.e., some measure of likelihood of that particular configuration. Thus, the potential between two neighboring deformation vectors  $\mathbf{d}_i$  and  $\mathbf{d}_j$  should be higher if both vectors are closer in direction and lower if they are more opposite in direction. Similarly, the potential between neighboring occlusion indicators  $\omega_i$  and  $\omega_j$  should be higher if their values are the same and lower if they are different. In order to simplify understanding and notation of the model, we collapse each pair of quantities  $(\mathbf{d}_i, \omega_i)$  into a single 3-D vector quantity  $\mathbf{h}_i = [\mathbf{d}_i, \omega_i]^T$ . We chose to use the following separable potential function between any pair of neighboring vectors  $\mathbf{h}_i$  and  $\mathbf{h}_j$ :

$$\Psi_{i,j}(\mathbf{h}_i, \mathbf{h}_j) = \Psi_{\mathbf{d},i,j}(\mathbf{d}_i, \mathbf{d}_j) \cdot \Psi_{\omega,i,j}(\omega_i, \omega_j) \quad (5)$$

where the component potential functions for neighboring

deformations and occlusions are given as follows:

$$\Psi_{\mathbf{d},i,j}(\mathbf{d}_i, \mathbf{d}_j) = \exp \left\{ -\frac{1}{2} \left( a \|\mathbf{d}_i\| + a \|\mathbf{d}_j\| + b \|\mathbf{d}_i - \mathbf{d}_j\| \right) \right\} \quad (6)$$

$$\Psi_{\omega,i,j}(\omega_i, \omega_j) = \begin{cases} \alpha_0, & \omega_i = \omega_j = 0 \\ \alpha_1, & \omega_i = \omega_j = 1 \\ \alpha_2, & \omega_i \neq \omega_j \end{cases} \quad (7)$$

Parameters  $a$  and  $b$  represent penalties on absolute and relative deformations, respectively, while parameters  $\alpha_i$  correspond to priors on their corresponding occlusion configuration  $(\omega_i, \omega_j)$ . We used  $a = 0.05$ ,  $b = 0.1$ ,  $\alpha_0 = 0.7$ ,  $\alpha_1 = 0.14$ , and  $\alpha_2 = 0.08$ . These parameters were learned from training data using the expectation-maximization (EM) algorithm [7] to maximize performance on a small subset of the CASIA database.

The potential  $\Psi_i(\mathbf{h}_i, \mathbf{O}_i)$  over each group  $\mathbf{d}_i$ ,  $\omega_i$ , and the corresponding observation  $\mathbf{O}_i$  should be higher when the values of  $\mathbf{d}_i$  and  $\omega_i$  agree more closely with the observation at that region. Thus, it should be a function of both the occlusion statistic  $\pi_i$  and the match score  $m_i(\mathbf{d}_i)$  at the given value of the shift vector  $\mathbf{d}_i$ . Furthermore, it should depend only on the occlusion statistic  $\pi_i$  if  $\omega_i = 1$  (i.e., the region is thought to be occluded) and only on the match score  $m_i(\mathbf{d}_i)$  if  $\omega_i = 0$  (i.e., the region is thought to be unoccluded).

Let the random variable  $s$  be drawn from the distribution of “true” match scores, i.e., match scores observed at the (unknown) true shifts for given unoccluded iris regions. Similarly, let the random variable  $\pi$  denote the distribution of “true” occlusion metrics. The distributions  $P(s)$  and  $P(\pi)$  are assumed to be normally distributed with mean and variance parameters  $\mu_s, \sigma_s^2, \mu_\pi$ , and  $\sigma_\pi^2$ , and we define

$$F_s(S) \triangleq P(s < S) = \int_{-\infty}^S \mathcal{N}(s; \mu_s, \sigma_s^2) ds$$

$$F_\pi(\Pi) \triangleq P(\pi < \Pi) = \int_{-\infty}^\Pi \mathcal{N}(\pi; \mu_\pi, \sigma_\pi^2) d\pi \quad (8)$$

to be the cumulative distribution functions (cdfs) of  $s$  and  $\pi$ . The parameters of these distributions are learned from training data. We compute the probabilities  $F_s(m(\mathbf{d}_i))$  and  $F_\pi(\pi_i)$  of having observed *at least* the true match score and occlusion metric for each pixel or region, respectively. We then achieve monotonic potential functions  $\Psi_i(\mathbf{h}_i, \mathbf{O}_i)$  by setting them equal to the corresponding probability for the believed state of  $\omega_i$  as follows:

$$\Psi_i(\mathbf{h}_i, \mathbf{O}_i) = \begin{cases} F_s(m(\mathbf{d}_i)), & \omega_i = 0 \\ F_\pi(\pi_i), & \omega_i = 1 \end{cases} \quad (9)$$

We used as parameter values  $\mu_s = 0.4$ ,  $\sigma_s = 0.1$ ,  $\mu_\pi = 0$ , and  $\sigma_\pi = 0.14$ . These values were estimated by using only the first three images of each class in the CASIA database as training data. Some of the regions in this training set were manually labeled with occlusion state and/or true deformation, and only these regions were used in the estimation.

### 3.4. Belief propagation

Given a particular set of observations for nodes  $\mathbf{O}_i$ , the structure in Fig. 4 reduces to a Markov random field (MRF) [10] with potential functions specified in Section 3.3. In order to compute an overall match score, we want to first estimate the conditional distributions  $P(\mathbf{h}_i | \mathbf{O})$  for  $i = 1, \dots, N_s$ , where  $\mathbf{O}$  denotes the set of all observations  $\mathbf{O}_1, \dots, \mathbf{O}_{N_s}$ . One way to do this is to use loopy belief propagation (LBP) [5],[6]. LBP is a form of iterative optimization over the joint distribution represented by a graphical model. While LBP is not guaranteed to converge, we have observed in our experiments that the algorithm typically converges in 3 iterations or fewer.

We used code from Kevin Murphy’s MATLAB toolbox [8] to run LBP on this MRF for each query-template comparison. In each iteration of this LBP implementation, each unobserved node “passes” a message to each of its unobserved neighbors. The message  $\delta_{j \rightarrow k}^i$  from node  $j$  to neighboring node  $k$  at iteration  $i$  is computed according to the following equation:

$$\delta_{j \rightarrow k}^i(\mathbf{h}_k) = \sum_{\mathbf{h}_j} \Psi_j(\mathbf{h}_j, \mathbf{O}_j) \Psi_{j,k}(\mathbf{h}_j, \mathbf{h}_k) \times \prod_{l \in \mathcal{N}(j)-k} \delta_{l \rightarrow j}^{i-1}(\mathbf{h}_j), \quad (10)$$

where  $\Psi_j$  and  $\Psi_{j,k}$  are defined in (9) and (7), respectively, and  $\mathcal{N}(j)$  denotes the set of all neighbors of node  $j$  in the graph (Fig. 4).

The belief  $\hat{P}(\mathbf{h}_j | \mathbf{O})$  of quantity  $\mathbf{h}_j$  after iteration  $i$  can be computed via the following equation:

$$\hat{P}(\mathbf{h}_j | \mathbf{O}) = \frac{1}{Z_j} \Psi_j(\mathbf{h}_j, \mathbf{O}_j) \prod_{k \in \mathcal{N}(j)} \delta_{k \rightarrow j}^i(\mathbf{h}_j) \quad (11)$$

where the normalizing constant  $Z_j$  is given by

$$Z_j = \sum_{\mathbf{h}_j} \Psi_j(\mathbf{h}_j, \mathbf{O}_j) \prod_{k \in \mathcal{N}(j)} \delta_{k \rightarrow j}^i(\mathbf{h}_j). \quad (12)$$

We can then compute the marginal beliefs for each region as follows:

$$\begin{aligned} \hat{P}(\mathbf{d}_j | \mathbf{O}) &= \sum_{\omega_j} \hat{P}(\mathbf{h}_j | \mathbf{O}); \\ \hat{P}(\omega_j | \mathbf{O}) &= \sum_{\mathbf{d}_j} \hat{P}(\mathbf{h}_j | \mathbf{O}). \end{aligned} \quad (13)$$

We allowed a maximum of 3 iterations for each run of LBP.

	CASIA			
	FAR = 1%	0.1%	0.01%	0.001%
Baseline	1.0%	2.3%	4.8%	9.6%
Proposed	0%	0.1%	0.7%	1.9%
	ICE Experiment 1			
	FAR = 1%	0.1%	0.01%	0.001%
Baseline	2.35%	3.45%	4.82%	6.26%
Proposed	0.17%	0.33%	0.82%	1.56%
	ICE Experiment 2			
	FAR = 1%	0.1%	0.01%	0.001%
Baseline	2.16%	3.02%	3.92%	5.25%
Proposed	0.64%	0.94%	1.26%	1.91%

Table 1. FRR achieved at four values of FAR on CASIA and ICE databases.

### 3.5. Score computation

Once the conditional distributions  $P(\mathbf{d}_i | \mathbf{O})$  have been estimated, we compute a single match score  $M_i$  for each subregion  $i$ . Specifically, for each subregion we compute the expected match score with respect to the estimated distribution as follows:

$$M_i = \sum_{\mathbf{d}} m_i(\mathbf{d}) \hat{P}(\mathbf{d}_i = \mathbf{d} | \mathbf{O}). \quad (14)$$

Note that in computing each subregion match score  $M_i$ , we want to measure only the similarity and not the occlusion, i.e., we do not want a high probability of occlusion to cause a low match score. For this reason, we use the marginal distribution  $\hat{P}(\mathbf{d}_i | \mathbf{O})$  rather than the joint distribution  $\hat{P}(\mathbf{h}_j | \mathbf{O})$ . We compute the overall comparison score  $M$  from the  $N_s$  individual subregion scores  $M_i$  as a normalized weighted sum of the subregion match scores:

$$M = \frac{\sum_{i=1}^{N_s} \beta_i M_i}{\sum_{i=1}^{N_s} \beta_i} \quad (15)$$

where the weights

$$\beta_i = \hat{P}(\omega_i = 0 | \mathbf{O}) \quad (16)$$

are equal to the beliefs of non-occlusion for each region.

## 4. Experimental setup and results

We implemented a baseline method to which to compare our proposed method. The baseline method does not use the deformation model described in Section 3.1; rather, it assumes that the deformation between template and query iris codes is a rigid shift of the entire iris code. The overall match score for each iris comparison is then computed

using (4); no belief propagation is done, and the equations in Section 3.5 are not used since we essentially have only one subregion. This method is close to that of Daugman reported in [3], which is considered to be a standard iris matching approach [13]; the difference is that the parameter values used for the Gabor wavelets may be different. Because the values used by Daugman were not specified, we chose values to optimize performance of the baseline method.

Verification experiments were carried out on both the CASIA and ICE databases. For each database, we computed a match score between every pair of images in the database, resulting in 1,014,600 total comparisons for ICE Experiment 1 (right-eye), 1,166,628 for ICE Experiment 2 (left-eye), and 285,390 for CASIA. For a potential match threshold, we can compute the percentage of within-class match scores below the threshold (called the false reject rate) and the percentage of between-class match scores exceeding the threshold (called the false accept rate). We then found the thresholds that resulted in false-accept rates (FARs) of 1%, 0.1%, 0.01%, and 0.001% and reported the false reject rates (FRRs) at these thresholds in Table 1.

Fig. 5 contains normalized histograms showing the separation of match scores between authentic and impostor comparisons on the ICE Experiment 1 database. At FAR = 0.1% (a typical ICE performance standard), we observe a reduction in FRR by a factor of more than 17 for CASIA and more than 10 for ICE Experiment 1. ICE Experiment 2 was known to have at least one mislabeled image, which likely caused an increased error rate.

Because the proposed algorithm is significantly more complex than the baseline algorithm, we observed that the computational load per comparison increased by a factor of approximately 4 in our implementation. On average, each baseline comparison took  $\sim 60$  ms. This value increased

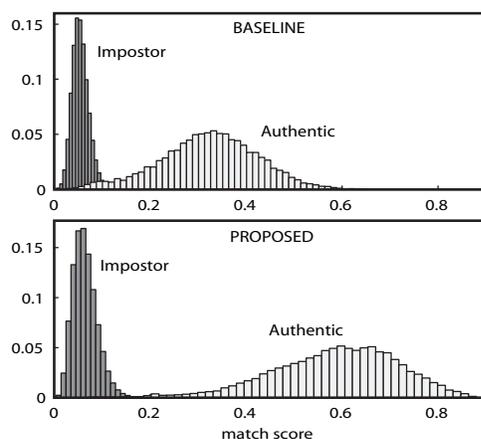


Figure 5. Histograms of match score distributions for both algorithms on the ICE Experiment 1 database.

to  $\sim 250$  ms in the proposed algorithm, of which  $\sim 150$  ms was devoted to computing the observed data in the graphical model and the remainder to running LBP and computing the final match score.

## 5. Conclusions

We have explained and demonstrated the use of an undirected lattice-type graphical model for modeling non-linear deformations and local occlusions in iris images. The main advantage of this model is the ability to compute distributions over local deformations rather than point estimates, thereby preserving information in the match score arrays. Also, by estimating occlusion and deformation jointly rather than separately, the observations of each can be combined probabilistically to yield better estimates of both. We showed that, by using this model, verification error rates on two iris databases can be reduced by more than an order of magnitude compared to a standard method.

## References

- [1] CASIA iris database. Institute of Automation, Chinese Academy of Sciences. <http://nlpr-web.ia.ac.cn/english/irds/irisdatabase.htm>, 2004. 2
- [2] ICE iris database. National Institute of Standards and Technology. <http://iris.nist.gov/ICE/>, 2005. 2
- [3] J. Daugman. High-confidence recognition of persons by iris patterns. In *Proc. IEEE 35th International Carnahan Conference on Security Technology*, pages 254–263, 2001. 1, 2, 3, 6
- [4] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. Wiley, 2001. 3
- [5] P. Felzenszwalb and D. Huttenlocher. Markov random fields and images. *Intl. Journal of Computer Vision*, 70(1):41–54, 2006. 5
- [6] B. J. Frey and N. Jojic. A comparison of algorithms for inference and learning in probabilistic graphical models. *IEEE Trans. Patt. Anal. Mach. Intell.*, 27(9):1392–1416, 2005. 1, 5
- [7] M. I. Jordan. *Learning in Graphical Models*. M.I.T. Press, Cambridge, MA, 1999. 4
- [8] K. Murphy. Conditional random fields (CRF) toolbox for MATLAB. <http://www.cs.ubc.ca/~murphyk/Software/CRF/crf.html>. 5
- [9] M. Pawan Kumar, P. H. S. Torr, and A. Zisserman. Learning layered motion segmentation of video. In *Proc. IEEE Intl. Conf. on Computer Vision (ICCV)*, pages 33–40, 2005. 1
- [10] P. Perez. Markov random fields and images. *CWI Quarterly*, 11(4):413–437, 1998. 5
- [11] J. Thornton, M. Savvides, and B. V. K. Vijaya Kumar. Enhanced iris matching using estimation of in-plane nonlinear deformations. In *Proc. of SPIE Defense and Security Symposium*, Apr. 2006. 1, 2
- [12] R. Wildes. Iris recognition. In J. Wayman, A. Jain, D. Maltoni, and D. Maio, editors, *Biometric Systems*, chapter 3, pages 63–95. Springer, 2005. 1
- [13] G. O. Williams. Iris recognition technology. *IEEE AES Systems Magazine*, pages 23–29, April 1997. 6
- [14] O. M. C. Williams, A. Blake, and R. Cipolla. The variational ising classifier (vic) algorithm for coherently contaminated data. In *Advances in Neural Information Processing Systems 17*, pages 1497–1504. 2005. 1
- [15] J. Winn and N. Jojic. Locus: Learning object classes with unsupervised segmentation. In *Proc. IEEE Intl. Conf. on Computer Vision (ICCV)*, 2005. 1