

A Mosaicing Scheme for Pose-Invariant Face Recognition

Richa Singh, *Student Member, IEEE*, Mayank Vatsa, *Student Member, IEEE*,
Arun Ross, *Member, IEEE*, and Afzel Noore, *Member, IEEE*

Abstract—Mosaicing entails the consolidation of information represented by multiple images through the application of a registration and blending procedure. We describe a face mosaicing scheme that generates a composite face image during enrollment based on the evidence provided by frontal and semiprofile face images of an individual. Face mosaicing obviates the need to store multiple face templates representing multiple poses of a user's face image. In the proposed scheme, the side profile images are aligned with the frontal image using a hierarchical registration algorithm that exploits neighborhood properties to determine the transformation relating the two images. Multiresolution splining is then used to blend the side profiles with the frontal image, thereby generating a composite face image of the user. A texture-based face recognition technique that is a slightly modified version of the C2 algorithm proposed by Serre *et al.* is used to compare a probe face image with the gallery face mosaic. Experiments conducted on three different databases indicate that face mosaicing, as described in this paper, offers significant benefits by accounting for the pose variations that are commonly observed in face images.

Index Terms—Face mosaicing, face recognition, multiresolution splines, mutual information.

I. INTRODUCTION

THE PROBLEM of 2-D face recognition continues to pose challenges even after several years of research in this field [1]. State-of-the-art algorithms exhibit various degrees of sensitivity to changes in illumination, pose, facial expressions, accessories, etc. Designing pose-invariant algorithms is particularly very challenging, as discussed in the Face Recognition Vendor Test 2002 report [2]. Several methods have been suggested to address the issue of pose variations including the use of active appearance models [3], morphable models [4], 3-D facial imaging [5], multiple templates [6], and multiclassifier fusion [7], [8]. In this paper, we propose an image fusion scheme to generate the 2-D face mosaic of an individual during enrollment that can be successfully used to match various poses of a person's face during authentication. Mosaicing uses the frontal and side-profile face images (2-D) of a user to generate an extended 2-D image. The goal is to adequately characterize

an individual's face in 2-D plane, without attempting to compute the 3-D structure of the face. This avoids the complexity of generating 3-D structure information from multiple registered 2-D images. Mosaicing also obviates the need to store multiple templates of a user during enrollment, thereby optimizing storage demands and processing time.

The potential of mosaicing facial images has not received extensive attention in the literature. Table I summarizes the face mosaicing techniques proposed by researchers. Yang *et al.* [11] propose an algorithm to create panoramic face mosaics. Their acquisition system consists of five cameras that simultaneously obtain five different views of a subject's face. In order to determine the corresponding points in multiple face views, the authors place ten colored markers on the face. Based on these control points, their algorithm uses a series of fast linear transformations on component images to generate a face mosaic. Finally, a local smoothing process is carried out to smooth the mosaiced image. Two different schemes are used to represent the panoramic image: 1) one in the spatial domain and 2) the other in the frequency domain. The frequency representation and spatial representation are observed to result in an identification accuracy of 97.46% and an accuracy of 93.21%, respectively, on a database of 12 individuals.

Liu and Chen [12] describe a face mosaicing technique that uses a statistical model to represent the mosaic. Given a sequence of face images captured under an orthographic camera model, each frame is unwrapped onto a certain portion of the surface of a sphere via a spherical projection. A minimization procedure using the Levenberg–Marquardt algorithm is employed to optimize the distance between an unwrapped image and the sphere. The statistical representational model consists of a mean image and a number of eigenimages. The novelty of this technique is given as follows: 1) the use of spherical projection, as opposed to cylindrical projection, which works better when there is head motion in both the horizontal and vertical directions, and 2) the computation of a representational model using both the mean image and the eigenimages rather than a single template image. Although the authors state that this method can be used for face recognition, no experimental results have been presented in this paper. In [13], the authors propose another algorithm in which the human head is approximated with a 3-D ellipsoidal model. The face, at a certain pose, is viewed as a 2-D projection of this 3-D ellipsoid. All 2-D face images of a subject are projected onto this ellipsoid via geometrical mapping to form a texture map that is represented by an array of local patches. Matching is accomplished by adopting a probabilistic model to compute

Manuscript received May 23, 2006; revised January 27, 2007. This paper was recommended by Guest Editor N. Ratha.

The authors are with the Lane Department of Computer Science and Engineering, West Virginia University, Morgantown, WV 26506-6109 USA (e-mail: richas@csee.wvu.edu; mayankv@csee.wvu.edu; arun.ross@mail.wvu.edu; afzel.noore@mail.wvu.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSMCB.2007.903537

TABLE I
COMPARISON OF THREE EXISTING FACE MOSAICING SCHEMES

	Yang et. al. [11]	Liu & Chen [12]	Singh et. al.[15]
Subjects	12	68	12 + 15
Registration	Affine using points	Affine using triangles	Affine using regions
Mosaicing	Concatenation	Geometrical mapping	Multi-resolution splines
Representation	Spatial: PCA, Frequency: FFT amplitude	Statistical model	Local binary pattern

the distance of patches from an input face image. The authors report an identification accuracy of 90% on the CMU Pose, Illumination, and Expression (PIE) database [14].

Face mosaicing has also been used in nonbiometric applications such as facial animation and rendering [16], and 3-D face image generation [17]. However, these algorithms generate the face mosaic using complex models that do not necessarily preserve the biometric features of the face.

The concept of mosaicing may be viewed as an exercise in *information fusion*. When multiple images of a subject's face are available at the time of enrollment, a common approach is to treat these images (also known as gallery images) as independent entities; thus, when a probe (query) image is presented to the system, it is compared against each gallery image *independently*, and the resulting set of scores consolidated to generate a single score (e.g., via the sum rule) indicating the proximity of the probe image with the subject in the database. This is fusion at the match-score level. However, in the case of mosaicing, multiple images of a subject's face are fused into a single entity in the image domain itself. Hence, this could be viewed as fusion at the raw-data (i.e., image) level.

The remainder of this paper is organized as follows. The proposed face mosaicing algorithm is described in Section II. In Section III, we present a face recognition algorithm for matching mosaiced faces. The proposed face recognition algorithm is a modification of the C2-feature-based face recognition algorithm [18], [19]. Section IV discusses the database used to validate the performance of the proposed face mosaicing and recognition algorithm. Experimental results are presented in Section V, and conclusions are presented in Section VI.

II. PROPOSED FACE MOSAICING ALGORITHM

This section describes the face mosaicing algorithm used to consolidate the evidence presented by multiple pose images of the same face. It is assumed, therefore, that, at the time of enrollment, multiple poses of an individual's face are available. The face is segmented (localized) from each image using the gradient vector flow technique (see [20] for details). A pair of face images, typically representing the frontal and profile views of an individual, are mosaiced after aligning them using a hierarchical registration algorithm. Registered images are mosaiced using the multiresolution splines algorithm based on Gaussian and Laplacian pyramids [21]. Multiresolution splines also perform blending as an integral part of mosaicing, thereby offering some inherent advantages.

A. Hierarchical Registration Model

Before mosaicing, it is necessary to transform the images obtained during enrollment into a common image domain. The process of finding the transform that aligns one image to another is called image registration. As described earlier in Section I, existing face mosaicing algorithms use some form of affine transformation for registration. However, these algorithms do not consider the nonlinear deformation that is present in the images. In this section, we propose a hierarchical registration algorithm in which we first perform approximate registration using an affine transformation model [22]. The affine transformed images are then finely registered using a mutual-information-based registration algorithm [23], [24], resulting in more exact alignment between the images.

1) *Affine Transformation Model*: Let $I_1 = I(x, y, t_1)$ and $I_2 = I(x, y, t_2)$ be the two images to be mosaiced. Here, I_1 is the source image, and I_2 is the target image, i.e., I_1 has to be suitably transformed in order to align it with I_2 . The pixel coordinates are represented using x and y spanning the domain of the image. The relationship of the pixels and their intensities between the two images can be modeled as

$$I(x, y, t_1) = \frac{I(a_1x + a_2y + a_3, a_4x + a_5y + a_6, t_2) - a_7}{a_8} \quad (1)$$

where $a_1, a_2, a_4,$ and a_5 are the affine parameters summarizing the rotation, scaling, and shear; a_3 and a_6 are the translation parameters; and a_7 and a_8 are the parameters that embody changes in brightness and contrast, respectively. The following error function is minimized in order to estimate these parameters:

$$\text{Error}(\mathbf{a}) = \sum_{x,y} \left[I(x, y, t_1) - \frac{I(a_1x + a_2y + a_3, a_4x + a_5y + a_6, t_2) - a_7}{a_8} \right]^2 \quad (2)$$

where $\mathbf{a} = (a_1, a_2, \dots, a_8)^T$. Approximating this error function using the first-order truncated Taylor series expansion gives [22]

$$\text{Error}(\mathbf{a}) = \sum_{x,y} \left[I_t(x, y, t) - \frac{(a_1x + a_2y + a_3 - x)I_x(x, y, t)}{a_8} + \frac{(a_4x + a_5y + a_6 - y)I_y(x, y, t) - a_7}{a_8} \right]^2 \quad (3)$$

where $I_x(\cdot)$, $I_y(\cdot)$, and $I_t(\cdot)$ are the spatial and temporal derivatives of $I(\cdot)$. This error function can be minimized by differentiating with respect to \mathbf{a} , i.e., $d\text{Error}(\mathbf{a})/d\mathbf{a} = 0$. The solution of (3) is

$$\mathbf{a} = \left[\sum_{x,y} A^T A \right]^{-1} \left[\sum_{x,y} A b \right] \quad (4)$$

where $A = (xI_x, yI_x, xI_y, yI_y, I_x, I_y, -I, -1)^T$, and $b = I_t - I + xI_x + yI_y$. The output of this minimization problem gives the eight optimal parameters. Using these parameters, the source face image I_1 is transformed to obtain the registered face image I_R . In this manner, the affine model performs a coarse registration of the two images. However, local regions such as the eyes and mouth still need to be finely aligned in order to remove the degeneracies that are present in the transformed image.

2) *Mutual-Information-Based Transformation Model*: In two coarsely registered face images, the neighborhood of the corresponding regions may not be rigorously identical. This is due to differences in the geometry and local deformations that are present in the constituent images. Subpixel shifting can also occur, leading to differences in the two images. Thus, a fine alignment is necessary to account for these nonlinear deformations.

Let I_2 be the target face image and I_R be the face image, which is coarsely registered with respect to I_2 using the affine transformation that is described in Section II-A1. For fine registration, we transform I_R such that the mutual information between I_2 and I_R is maximized [24]. Mutual-information-based image registration is widely used in medical imaging [24] and other related applications. Mutual information between two face images $M(I_2, I_R)$ can be represented as

$$M(I_2, I_R) = H(I_2) + H(I_R) - H(I_2, I_R) \quad (5)$$

where $H(\cdot)$ is the entropy of the image. Here, maximizing the mutual information involves maximizing the entropy, i.e., $H(I_2)$ and $H(I_R)$, and minimizing the joint entropy $H(I_2, I_R)$. Two images are considered to be optimally registered when the mutual information between them is maximum. Mutual-information-based registration is however sensitive to changes that occur in the distributions, as a result of differences in overlap region. So, we use the weighted normalized mutual information for face images, which can be written as

$$\hat{M}(I_2, I_R, c) = c \left(\frac{H(I_2) + H(I_R)}{H(I_2, I_R)} \right) \quad (6)$$

where c ($0 < c \leq 1$) is a weighting parameter that controls the amount of localization in the similarity measure. This method is a modification over the one proposed by Hill *et al.* [23]. The modified function that is represented by (6) is used for the fine registration of I_2 and I_R . The transformation parameters computed from the affine transformation model \mathbf{a} are used




Fig. 1. Image registration using the proposed hierarchical registration algorithm. Frontal and profile images are first placed at the center of a 256×256 image space. (a) Input profile image. (b) Input frontal image. (c) Profile image registered with respect to the target frontal image.




Fig. 2. Masks generated from two profile images.

as the initial transformation parameters, along with $c = 0.01$. Thus, a total of nine parameters are used as the initial transformation parameters for this model. A set of mutual information values between I_2 and I_R is computed by varying the parameters with a small step size (≈ 0.01) in both the positive and negative directions. Parameters are varied in the range of $[-\mathbf{a}/5, \mathbf{a}/5]$. The mutual information is computed for all possible combinations of the parameters, and the transformation parameters corresponding to the maximum mutual information are selected for fine registration. To account for the nonlinearity that is present in face images, the mutual-information-based registration is performed in blocks with a size of 8×8 . The coarsely registered profile image I_R is transformed using these parameters to obtain the final registered image. Fig. 1 shows a profile image I_1 transformed with respect to the frontal image I_2 .

B. Mask Generation

Once a pair of images are registered, the next step involves blending the two images into a single entity. This entails the development of a spatial mask indicating the pixelwise contribution of the individual images to the final mosaiced image. Since the facial structure is different for every individual, a dynamic runtime mask generation algorithm is used. The mask is computed using the local phase correlation between the two images. The two images are first tessellated into blocks of size 8×8 . Next, the phase correlation between corresponding blocks from the two face images is computed. When a correlation peak is observed, the block is labeled as a *match*; otherwise, it is labeled as a *nonmatch*. This results in a cluster of match/nonmatch blocks. The boundary of the matched *region* is selected as the boundary of the mask.




Fig. 3. Levels in the Gaussian pyramid expanded to the original size to see the effects of the low-pass filter. (a) Level 0, (b) level 1, (c) level 2, and (d) level 3.

The mask values on one side of the boundary is set to “0,” while the other side is set to “1” (e.g., “1” may correspond to the frontal face image and “0” to the profile face image). If there are any isolated blocks with the label *match*, they are reassigned the label *nonmatch* and do not contribute to the mask boundary. Generally, the correlation peak is found in a thin vertical region containing the eye (left eye for left profile image and right eye for right profile image). Fig. 2 shows the sample masks that are generated for a left and right profile face image with respect to an arbitrary frontal image. This mask is used during the blending process as will be seen in the next section.

C. Stitching and Blending

For blending the two images into a single mosaic, we use multiresolution splines [21]. Image splining (i.e., blending) can be performed based on a simple spline-weighting function¹ straddling the boundary of the two images, but the quality of the stitched image depends on the step size (or window) that is chosen. A large step size may lead to blurring, whereas a small step size may result in discontinuities at the boundary. To overcome this problem, Burt and Adelson [21] used multiresolution splines to determine different step sizes for the various frequency components constituting the boundary. The crux of this technique involves computing a Gaussian pyramid of subimages, followed by a Laplacian pyramid, based on the two images to be mosaiced; the pyramid structure is used to estimate the spline weighting function that relies on the frequency-domain information of the image.

A sequence of low-pass-filtered images is obtained by iteratively convolving each of the constituent images with a 2-D Gaussian filter kernel. The resolution and sample density of the image between successive iterations (levels) is reduced, and therefore, the Gaussian kernel operates on a reduced version of the original image in every iteration. The resultant images G_0, G_1, \dots, G_N may be viewed as a “pyramid,” with G_0 having the highest resolution (lowermost level) and G_N having the lowest resolution (uppermost level) (see Fig. 3). Let $w(m, n)$ represent the Gaussian kernel with dimensions of 5×5 and a reduction factor of 4. The *reduce* operation can be written as

$$\text{Reduce}(I(i, j)) = \sum_{m=1}^5 \sum_{n=1}^5 w(m, n) I(2j + m, 2j + n). \quad (7)$$

¹The term splining is used to refer to a transition function that indicates the weighting of pixels that are associated with the two images at the boundary (see [21] for details).




Fig. 4. Laplacian pyramid of a profile image from level 0 to level 6.



Fig. 5. Laplacian pyramid of a frontal image from level 0 to level 6.

A Gaussian pyramid G_l is defined as

$$G_0 = I \quad (8)$$

$$G_l = \text{Reduce}[G_{l-1}], \quad 0 < l < N. \quad (9)$$

As shown in Fig. 3, the effect of convolution is to blur the image, thereby reducing the filter band limit by an octave between levels while reducing the sample density by the same factor. The Gaussian pyramid has the effect of a low-pass filter to soften the edges of the mask.

The multiresolution spline, as described in [21], requires bandpass images, as opposed to low-pass images. Bandpass images are computed by interpolating (resizing) the image at each level of the Gaussian pyramid and then subtracting it from the next lowest level. This results in a sequence of bandpass images that may be viewed as a Laplacian pyramid (L_0, L_1, \dots, L_N), as shown in Figs. 4 and 5. The term Laplacian is used since the Laplacian operator resembles the difference of Gaussian-like functions. These bandpass images are a result of convolving the difference of two Gaussians with the original image. The steps used to construct this pyramid can also be used to exactly recover the original image.



Fig. 6. Laplacian pyramid of the mosaiced image from level 0 to level 6.

The process described previously may be summarized as follows:

$$L_l = G_l - \text{Expand}[G_{l+1}], \quad 0 \leq l < N. \quad (10)$$

Here, the $\text{Expand}[\cdot]$ operator interpolates a low-resolution image to the next highest level and can be written as

$$G_{l,k}(i, j) = 4 \sum_{m=-2}^2 \sum_{n=-2}^2 w(m, n) G_{l,k-1} \left(\frac{i-m}{2}, \frac{j-n}{2} \right). \quad (11)$$

Note that $G_{l,k}$ in (11) denotes “expanding” G_l k number of times. Various features of the face are segregated by scale in different levels of the pyramid. Hence, as shown in Figs. 4–6, the textural features of face are preserved over multiple levels of the pyramid. Let L_1 and L_2 represent the Laplacian pyramids of the two images that are being splined (i.e., blended). Let GR be the pyramid associated with the Gaussian-weighted mask discussed in Section II-B. The multiresolution spline LS is then computed as

$$LS_l(i, j) = GR_l(i, j)L_{1l}(i, j) + (1 - GR_l(i, j))L_{2l}(i, j) \quad (12)$$

where l is the level of the pyramid. The splined images at various levels are expanded and summed together to obtain the final face mosaic, as shown in Fig. 7. Gradient-vector-flow-based active contour model is used to extract the face boundary of the mosaiced face.

III. FACE RECOGNITION USING MODIFIED C2 FEATURES

Several different texture-based face recognition algorithms have been proposed in the literature [25], [26]. However, most of the current texture-based face recognition algorithms fail to explicitly account for important spatial statistics between texture elements. Spatial statistics are important when analyzing facial textures that have similar texture frequency but differ in the distribution of texture elements. Face recognition algorithms should also handle small distortions and preserve the local feature geometry. In [18], a generic model-based feature extraction algorithm, which extracts visual features from an object using the fundamentals of a biological visual system, is proposed. The model is a feedforward hierarchy consisting of four layers of computational units: simple S units and complex C

units. The simple S units combine their inputs with Gaussian-like tuning to increase object selectivity. The complex C units pool their inputs through a maximum operation to achieve scale and shift invariance. The feature extractor is a Gabor filter bank consisting of 16 different filters and four different orientations, resulting in 64 different maps when applied to a circular image patch. Each filter is parameterized based on its scale, width, and frequency. The 64 maps are arranged in eight bands (see [18] for details about the filter bank). The C2 algorithm operates as follows.

- The input face image is subjected to the Gabor filter bank described previously, resulting in 64 maps arranged in eight bands. These maps constitute the S1 feature set.
- The C1 feature set is obtained by computing the maximum response over all scales in each band for all four orientations. A large pool of patches $P_{i=1, \dots, K}$ is extracted at random positions from all the training images. These patches are extracted for all four orientations, and the radii of these patches vary from 4 to 16.
- For each feature set C1, the value Y is computed across all the bands for all image patches X at all positions P using the following equation:

$$Y = \exp(-\gamma \|X - P\|^2) \quad (13)$$

where γ is the aspect ratio. These patches are set as the prototypes or centers of S2 units. The S2 units behave as radial basis function during recognition.

- The maximum over all positions and scales for each S2 map gives the C2 features, which are shift and scale invariant.

Serre *et al.* [18] have demonstrated superior recognition performance due to these features on different face and texture databases. For our purposes, we have modified this basic feature extraction algorithm to make it compatible with a face mosaicing application. In the modified algorithm, the filter bank is generated with 2-D log polar Gabor transforms. The functional form of a 2-D log polar Gabor G can be written as

$$G_{r_0, \theta_0}(r, \theta) = \exp(-2\pi^2 \sigma^2) \times \left[(\ln(r) - \ln(r_0))^2 s^2 + (\ln(r) \sin(\theta - \theta_0)) \right] \quad (14)$$

and the position of the filter in the Fourier domain is defined by

$$r_{00} = \sqrt{2}, \quad r_{0i} = 2^i * r_{00}, \quad \theta_{0i} = i * \frac{2\pi}{N_\theta} \quad (15)$$

where r_{00} is the smallest possible frequency, N_θ is the number of filters on the unit circle, and at index E , σ_E and s_E are further defined as

$$\sigma_E = \frac{1}{\ln(r_0)\pi \sin(\pi/N_\theta)} \sqrt{\frac{\ln 2}{2}} \quad (16)$$

$$s_E = \frac{\ln(r_0)\pi \sin(\pi/N_\theta)}{\ln 2} \sqrt{\frac{\ln 2}{2}}. \quad (17)$$




Fig. 7. (a)–(c) Frontal and profile input images. (d) Mosaiced face generated using (a)–(c). (e) Final mosaiced and cropped face image.

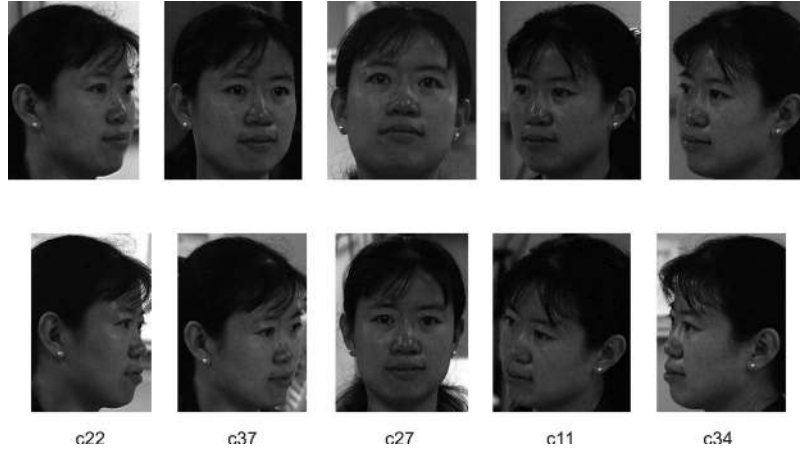


Fig. 8. Examples of images from the CMU PIE database.

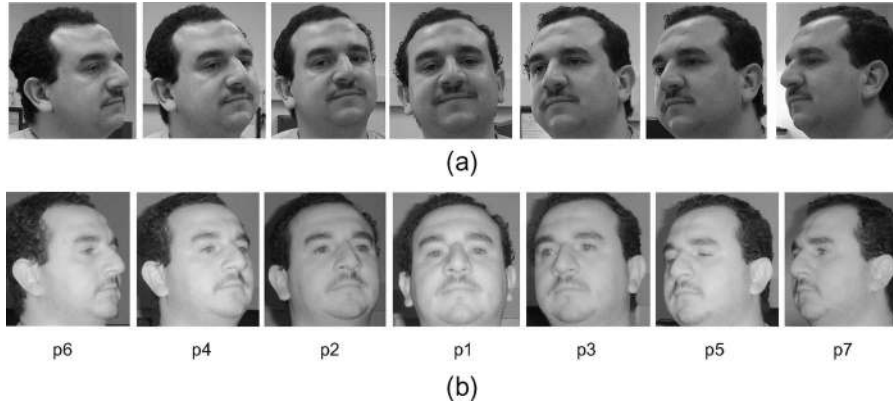


Fig. 9. Images from the WVU multispectral face database. (a) Visible spectrum. (b) SWIR spectrum.

According to Field [27], log Gabor functions are able to encode the images more efficiently compared to Gabor transforms since a Gabor transform would overrepresent the low-frequency components and underrepresent the high-frequency components. In addition, log polar Gabor functions impart rotation- and scale-invariant properties. In [18], 64 filters are used in the filter bank in order to mimic the functional ability of biological cells in the visual cortex. We modify this filter bank to include the Gaussian low-pass information. The filter bank therefore comprises of 16 filters at four orientations, along with the eight centers that surround differences of Gaussian filters and four low-pass Gaussian filters. So, the modified filter bank consists of 76 filters, which are used to compute the final C2 features. The size of the C2 feature vector varies for different face images. To efficiently match the mosaiced image with a non-mosaiced image, we use a learning-based 2ν support vector

machine (SVM) classifier [28], [29]. The classifier is trained with the features extracted from 148 mosaiced and 900 non-mosaiced face images of 108 individuals. This trained classifier is used to perform classification [30].

IV. DATABASE USED FOR EVALUATION

To validate the performance of the proposed face mosaicing and recognition algorithms, we used two databases.

- 1) CMU PIE Face Database: The CMU PIE [14] face database contains images from 68 individuals with variations in pose, illumination, and expression. For mosaicing, we use the images with variation in pose, i.e., the images labeled as c27, c37, c22, c11, and c34 from both sessions. Fig. 8 shows the two sets of images of an individual

TABLE II
IDENTIFICATION ACCURACY OF MULTIPLE FACE RECOGNITION
ALGORITHMS ON THE THREE DATABASES

Algorithm	CMU PIE Face Database [14]	WVU Visible-light Face Database	WVU SWIR Face Database
PCA [36]	53.81%	49.63%	52.94%
FLDA [37]	58.24%	54.19%	58.38%
LFA [38]	72.59%	71.37%	73.86%
2D log Gabor [39]	95.32%	96.22%	97.19%
Original C2 feature [18]	95.41%	96.24%	97.21%
Modified C2 feature	96.76%	97.06%	97.92%

from the CMU PIE database. For each individual, the set of face images with neutral expression is used as the gallery image set. The images from the three other sessions, which have minimum variation in lighting and expression, are chosen as probe images.

- 2) WVU Multispectral Face Database: We assembled a multispectral face database of 40 individuals. The database contains face images in both the visible and short-wave infrared (SWIR) light spectrum. For each spectral channel, images are captured in two different sessions, with seven images in each session. The seven images correspond to one frontal image, three left profile images, and three right profile images. Successive images are separated by a pose angle of approximately 20° . Fig. 9 shows examples of images from the database. Various researchers have demonstrated the superior performance of face recognition on SWIR images [31]–[35].


V. EXPERIMENTAL RESULTS

A. Performance Evaluation of Modified C2-Feature-Based Face Recognition Algorithm


We first validate the performance of the proposed face recognition algorithm described in Section III using the non-mosaiced face images. The 2ν -SVM classifier [28], [29] is trained using the C2 features extracted from the 620 labeled non-mosaiced training face images of the CMU PIE (68×5) and the WVU visible-light (40×7) databases. The classifier learns the genuine and impostor features from these training images prior to performing classification [18].

Recognition performance is evaluated separately for the two visible-light face databases. For evaluating the performance on the CMU PIE data set, all the training images are used as gallery images, and face images from the other sessions are used as probe images. Similarly, for the WVU visible-light face database, the training images are used as the gallery images, and the rest of the images are used as probe images to compute the identification performance. The performance of the proposed recognition algorithm is also compared with five existing face recognition algorithms.

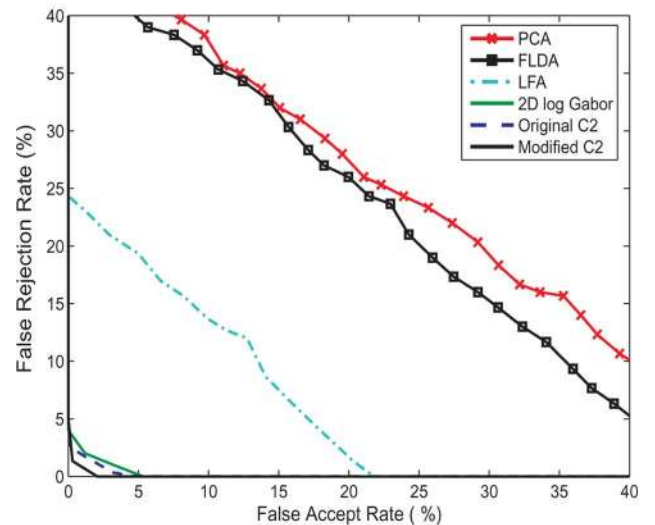
- 1) Principal component analysis (PCA) [36].
- 2) Fisher linear discriminant analysis (FLDA) [37].



(a)



(b)



(c)

Fig. 10. ROC indicating the performance of multiple face recognition algorithms on the three databases. (a) CMU PIE, (b) WVU visible-light, and (c) WVU SWIR face databases.

TABLE III
PERFORMANCE OF FACE MOSAICING BASED ON DIFFERENT INPUT IMAGE SEQUENCES

Database	Image Sequence	Verification Accuracy (0.001% FAR)	Identification Accuracy
CMU PIE Face Database [14]	c27	61.36%	61.41%
	c37 - c27 - c11	96.54%	96.88%
	c22 - c27 - c34	77.54%	75.98%
	c22 - c37 - c27 - c11 - c34	96.37%	96.40%
WVU Visible Face Database	p1	63.05%	64.51%
	p2 - p1 - p3	96.27%	96.35%
	p4 - p1 - p5	96.98%	97.52%
	p6 - p1 - p7	73.30%	71.19%
	p4 - p2 - p1 - p3 - p5	96.20%	96.37%
	p6 - p2 - p1 - p3 - p7	93.65%	93.44%
	p6 - p4 - p1 - p5 - p7	94.32%	94.35%
	p6 - p4 - p2 - p1 - p3 - p5 - p7	95.82%	95.80%
WVU SWIR Face Database	p1	64.47%	64.67%
	p2 - p1 - p3	97.13%	96.28%
	p4 - p1 - p5	97.91%	98.16%
	p6 - p1 - p7	74.02%	73.76%
	p4 - p2 - p1 - p3 - p5	96.99%	97.08%
	p6 - p2 - p1 - p3 - p7	96.26%	96.43%
	p6 - p4 - p1 - p5 - p7	95.84%	95.89%
	p6 - p4 - p2 - p1 - p3 - p5 - p7	94.70%	94.81%

- 3) Local feature analysis (LFA) [38] refers to a class of algorithms that extract a set of geometrical features and distances from facial images and use these features as the basis for face representation and comparison. Hopfield network is used to extract the output, which represents uncorrelated local features.
- 4) 2-D log polar Gabor transform [39].
- 5) Original C2 features [18].

A similar experiment is conducted for evaluating the performance on the SWIR database. For training the classifier, 280 (= 40 × 7) labeled non-mosaiced images from the WVU SWIR face database are used. These images are further used as gallery images, and the remaining 280 SWIR images are used as probe images. Identification accuracy for all the recognition algorithms is shown in Table II, and the receiver operating characteristic (ROC) plots are shown in Fig. 10. The results indicate that the modified C2-based feature extraction algorithm outperforms the other recognition algorithms for both the visible-light and SWIR face databases. The images contain variation in pose and expression, which leads to changes in the number and position of the features that are present in individual images. Therefore, several face recognition algorithms that are appearance-based or feature-based do not perform well. However, the texture-based recognition algorithm results in better performance and exhibits tolerance to these variations. Table II and the ROC plot in Fig. 10 show that the modified C2-feature-based face recognition algorithm gives an improvement

of 1.3% on the CMU PIE database and about 0.7% on the WVU multispectral database.

A mosaiced face image is expected to contain all the features in a face, while the frontal and side profile images have only a limited number of features. To match a mosaiced face with a non-mosaiced face image, we need a recognition algorithm that efficiently extracts the local features from the face and compares them by assigning proper weights to the features while appropriately accounting for missing features. To facilitate this, we use a local representation based on the textural features obtained using the modified C2 feature extraction algorithm in order to evaluate the performance of the proposed face mosaicing algorithm.

B. Performance Evaluation of the Face Mosaicing and Recognition Algorithm

The CMU PIE training set contains 68 mosaiced and 340 non-mosaiced labeled images, the WVU visible-light training set contains 40 mosaiced and 280 non-mosaiced labeled face images, and the WVU SWIR training set contains 40 mosaiced and 280 non-mosaiced labeled face images. The experimental setup here is used to compute the matching performance on all three databases separately.

- C2 features extracted from the mosaiced and non-mosaiced training images are used to train the 2ν-SVM-based classifier.




Fig. 11. Mosaiced images generated with the optimal input sequence. (a) CMU PIE face database [14]. (b) WVU visible-light face database. (c) WVU SWIR face database.

- C2 features extracted from the mosaiced face images are used as the gallery features.
- On all three databases, two sets of experiments are conducted.
 - 1) The first set of experiments is performed to compute the optimal sequence and number of face images that are used for mosaicing (Section V-B1).
 - 2) In the second experiment, the performance of match score fusion techniques [8] such as sum rule and min/max rules is compared with mosaiced images (image-level fusion) (Section V-B2).
 - 1) *Optimal Sequence for Face Mosaicing:* The CMU PIE database contains two left and two right profile face images with successive images having a difference of approximately

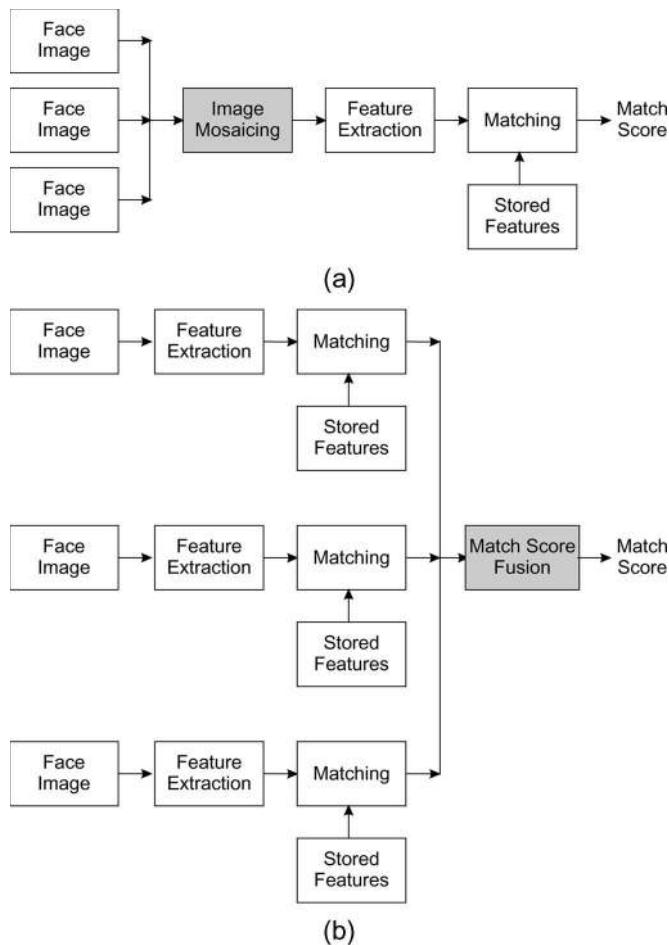


Fig. 12. Block diagram illustrating the difference between (a) image mosaicing and (b) match score fusion.

22.5°, while the WVU multispectral database contains three left and three right profile face images with a difference of about 20° between successive images. To find the optimal sequence and number of images for mosaicing, the frontal image of a subject is mosaiced with various combinations of the profile face images. In all sequences, an equal number of left and right profile images is used, and the difference in pose between successive profile images and the frontal image is approximately the same. This ensures that the extent of information on both “sides” of the mosaiced face is the same. Thus, we have four possible input sequences for the CMU PIE database [14] and eight possible input sequences for the WVU multispectral database.

Mosaiced face images are generated, with the training data set representing all possible input sequences. The verification and identification performance is evaluated by matching a mosaiced face image with all the probe non-mosaiced face images. Verification accuracy is computed at 0.001% false accept rate. Table III shows the results of this experiment. For the CMU PIE database, the input sequence “c37–c27–c11” results in the best verification and identification accuracy of 96.54% and 96.88%, respectively. For both visible-light and SWIR images from the WVU multispectral database, the input sequence “p4–p1–p5” results in the best performance. This suggests that increasing the number of input images for face mosaicing does not guar-

antee better performance. Experiments on all the databases indicate that the best result is obtained when the mosaic is generated using profile images having a difference of about 45° with the frontal image. Fig. 11 shows samples of mosaiced images that are generated using the optimal sequence for all three databases.

2) *Comparing Face Mosaicing With Match Score Fusion Algorithms:* We next compare the performance of fusion at the image level (mosaicing) with a few fusion operators at the match-score level (sum rule and min/max rules). Fig. 12 shows the steps that are involved in generating the final match score using these two methods. In the case of mosaicing, the gallery is assumed to contain only those mosaiced images corresponding to the optimal sequence obtained from the previous experiment. However, in the case of score fusion, all training images are used as gallery images. The experimental setup is summarized below.

- **Image fusion:** The gallery contains mosaiced images corresponding to the optimal sequence from one session, and the probe consists of the mosaiced and non-mosaiced images from the rest of the sessions. Images are matched using the proposed face recognition algorithm. Results of this experiment for the three databases are shown in the first row of Tables IV–VI.
- **Score fusion:** The gallery consists of all the non-mosaiced images from one session. The probe images correspond to all the non-mosaiced as well as mosaiced images (generated using the optimal input sequence) from the other sessions. Results are shown in the second and third rows of Tables IV–VI.

Tables IV–VI show the identification accuracies for all three experiments on CMU PIE, WVU visible-light, and WVU SWIR face databases, respectively. The results from the three tables are summarized below.

- The modified C2-feature-based face recognition method gives good performance for images having variations in size and content, making it particularly useful for recognizing mosaiced faces.
- A matching accuracy of 100% is obtained when both the gallery and probe images are mosaiced images.
- Matching mosaiced images with non-mosaiced images gives higher accuracy compared to matching a non-mosaiced image with other non-mosaiced profiles.
- The proposed face mosaicing algorithm (image-level fusion) gives better performance compared to the sum rule and min/max rules (score fusion algorithms).
- SWIR face recognition (using non-mosaiced as well as mosaiced images) results in better matching performance compared to visible-light face recognition. This may be attributed to the illumination-invariant characteristic of SWIR images. Fig. 13 illustrates the match scores generated using visible-light and SWIR images. The match scores are obtained using the modified C2-feature-based recognition algorithm. It is evident from this example

TABLE IV
IDENTIFICATION ACCURACY ON THE CMU PIE DATABASE BEFORE AND AFTER FUSION [14]

Gallery Image	Test Image					
	Mosaiced (%)	c27 (%)	c37 (%)	c22 (%)	c11 (%)	c34 (%)
Mosaiced	100	99.93	98.01	97.84	98.90	97.34
Sum Rule [8]	93.16	82.41	81.91	77.64	80.52	77.92
Min/Max Rule [8]	92.88	91.47	89.72	87.81	89.60	87.04
c27	98.91	99.89	84.07	80.76	83.96	79.16
c37	97.17	84.17	96.49	90.28	71.15	68.81
c22	97.02	80.71	90.01	95.61	66.84	62.18
c11	97.84	83.99	71.63	67.90	96.71	90.47
c34	97.07	79.10	69.47	63.37	90.11	95.90

TABLE V
IDENTIFICATION ACCURACY ON THE WVU VISIBLE-LIGHT FACE DATABASE BEFORE AND AFTER FUSION

Gallery Image	Test Image							
	Mosaiced	p1 (%)	p2 (%)	p4 (%)	p6 (%)	p3 (%)	p5 (%)	p7 (%)
Mosaiced	100	99.96	98.24	97.11	96.98	98.19	97.16	96.81
Sum Rule [8]	94.20	83.03	82.89	80.00	76.90	83.01	79.71	76.90
Min/Max Rule [8]	94.93	91.56	89.99	88.04	87.61	89.75	87.11	86.63
p1	99.92	99.95	89.16	85.18	81.70	89.53	84.92	80.65
p2	98.19	89.23	97.25	91.67	89.53	81.21	73.13	68.90
p4	97.13	85.26	91.93	97.01	91.16	75.33	68.46	63.08
p6	97.01	81.51	89.60	90.97	95.59	70.01	65.22	61.91
p3	98.07	89.59	80.76	74.51	70.31	97.34	91.53	87.86
p5	97.24	85.08	72.49	68.19	65.19	91.60	96.78	90.19
p7	96.90	81.13	68.31	62.87	62.00	88.17	90.47	95.29

TABLE VI
IDENTIFICATION ACCURACY ON THE WVU SWIR FACE DATABASE BEFORE AND AFTER FUSION

Gallery Image	Test Image							
	Mosaiced	p1 (%)	p2 (%)	p4 (%)	p6 (%)	p3 (%)	p5 (%)	p7 (%)
Mosaiced	100	100	98.79	98.11	97.36	98.70	98.12	97.22
Sum Rule [8]	95.08	90.86	90.03	87.64	83.96	90.81	85.69	82.71
Min/Max Rule [8]	95.73	93.41	92.71	90.51	89.32	92.24	90.06	88.83
p1	100	100	92.47	90.31	86.90	91.64	90.14	86.37
p2	98.71	91.49	98.75	93.22	90.51	91.29	89.73	85.98
p4	97.94	89.64	93.16	98.09	92.02	76.46	69.20	63.46
p6	97.07	86.58	90.37	91.88	96.99	71.24	65.79	62.90
p3	98.59	91.31	81.21	76.06	71.16	98.70	92.24	88.89
p5	98.00	89.67	73.19	69.20	65.87	92.26	97.51	91.25
p7	97.21	86.29	68.92	63.31	62.76	89.09	91.37	96.43

that the match scores obtained using SWIR images are more discriminating compared to that of visible-light face images.

- The performance of the proposed face mosaicing and recognition algorithm lies between 96.85% and 100% for all three databases.

Fig. 14 shows the ROC curve as a result of comparing mosaiced face images with other mosaiced and non-mosaiced images. The verification time using a mosaiced image is ~6 s,

while that using the score-fusion-based approach is between 8–12 s (in a Matlab environment). For the score fusion scheme, the matcher has to be invoked multiple times corresponding to each subject. Furthermore, if M is the memory required to store one face image (in bytes), the memory requirement without mosaicing is between $5M$ and $7M$, whereas it is approximately $1.1M$ for the mosaiced image. These results show that image mosaicing enhances the performance of face recognition algorithms while reducing the memory requirement and the matching time.

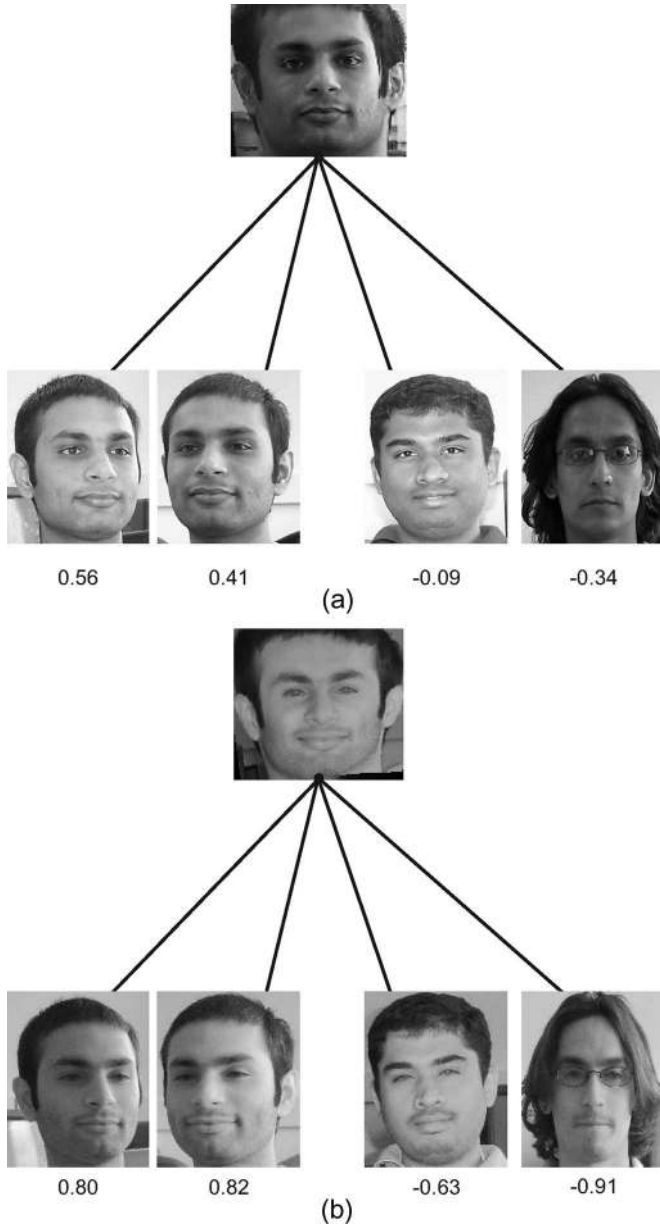
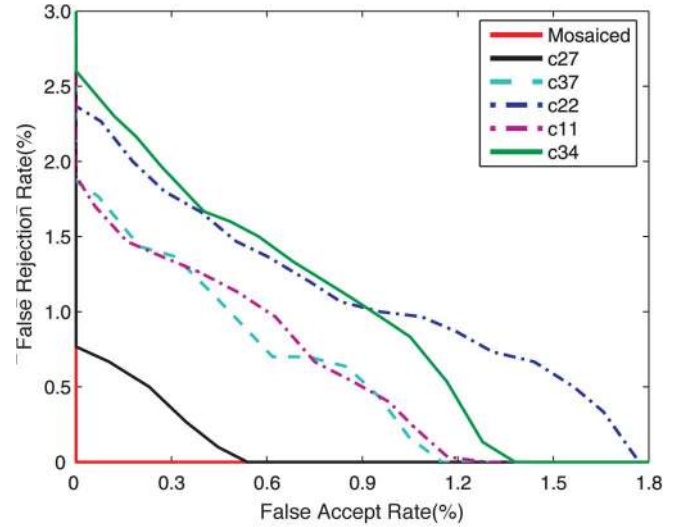


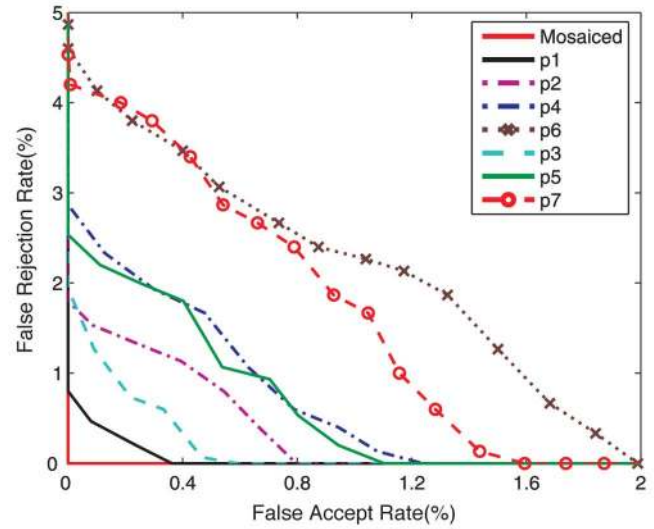
Fig. 13. Match scores obtained using the modified C2-feature-based algorithm. A value of “1” indicates a perfect match, while a “-1” represents a perfect reject. (a) WVU visible-light database. (b) WVU SWIR database.

VI. CONCLUSION

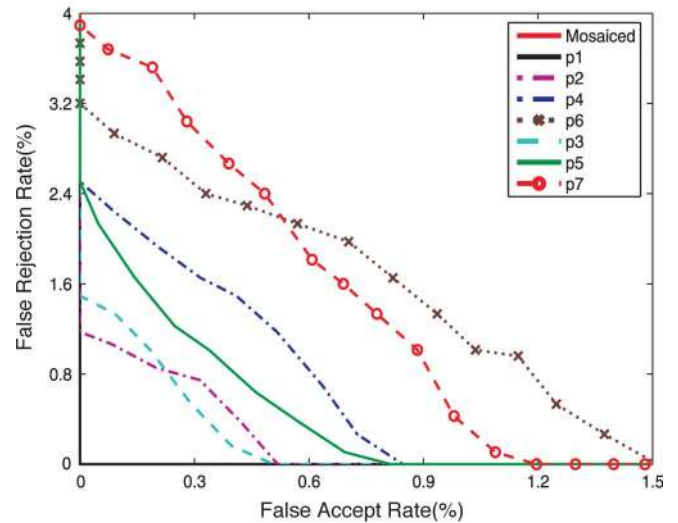
The primary goal of this paper is to demonstrate the role of face mosaicing in enhancing the matching performance of a face recognition system. Given multiple images of a face during enrollment, the mosaicing algorithm blends them into a single entity by employing a multiresolution splining scheme. Experiments reported on three different databases suggest that fusing information at the image level is better than fusing information at the match-score level in the context of this work. A modified version of the C2 algorithm, originally developed by Serre *et al.*, was used to extract features from the mosaiced and non-mosaiced face images. The modified algorithm is observed to perform very well both in the verification and identification modes of operation.



(a)



(b)



(c)

Fig. 14. Performance of the recognition algorithm with mosaiced face image as the gallery image (a) CMU PIE database [14]. (b) WVU visible-light database. (c) WVU SWIR database.

As part of our future work, we plan to evaluate the performance of 2-D face mosaicing against that of 3-D face models with respect to matching accuracy. Furthermore, we will investigate the role of mosaicing in perturbing the biometric content of the human face. Finally, we will examine the use of facial symmetry to create face mosaics when the frontal image is not available and only the left and right profiles are available.

ACKNOWLEDGMENT

The authors would like to thank R. Gross for granting access to the CMU PIE face database and the anonymous reviewers for their valuable comments.

REFERENCES

- [1] W. Zhao, R. Chellappa, A. Rosenfeld, and P. J. Phillips, "Face recognition: A literature survey," *ACM Comput. Surv.*, vol. 35, no. 4, pp. 399–458, 2003.
- [2] P. J. Phillips, P. Grother, R. J. Micheals, D. M. Blackburn, E. Tabassi, and J. M. Bone, *FRVT 2002: Evaluation Report*, 2003.
- [3] H. Kang, T. F. Cootes, and C. J. Taylor, "A comparison of face verification algorithms using appearance models," in *Proc. Brit. Mach. Vis. Conf.*, 2002, vol. 2, pp. 477–486.
- [4] V. Blanz, S. Romdhani, and T. Vetter, "Face identification across different poses and illuminations with a 3-D morphable model," in *Proc. Int. Conf. Autom. Face and Gesture Recog.*, May 2002, pp. 202–207.
- [5] K. I. Chang, K. W. Bowyer, and P. J. Flynn, "An evaluation of multimodal 2D+3D face biometrics," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 4, pp. 619–624, Apr. 2005.
- [6] U. Uludag, A. Ross, and A. K. Jain, "Biometric template selection and update: A case study in fingerprints," *Pattern Recognit.*, vol. 37, no. 7, pp. 1533–1542, 2004.
- [7] X. Lu, Y. Wang, and A. K. Jain, "Combining classifiers for face recognition," in *Proc. IEEE Int. Conf. Multimedia & Expo.*, Jul. 2003, vol. 3, pp. 13–16.
- [8] A. Ross and A. K. Jain, "Information fusion in biometrics," *Pattern Recognit. Lett.*, vol. 24, no. 13, pp. 2115–2125, 2003.
- [9] N. K. Ratha, J. H. Connell, and R. M. Bolle, "Image mosaicing for rolled fingerprint construction," in *Proc. Int. Conf. Pattern Recog.*, Aug. 1998, vol. 2, pp. 1651–1653.
- [10] A. Jain and A. Ross, "Fingerprint mosaicking," in *Proc. Int. Conf. Acoust., Speech, and Signal Process.*, May 2002, vol. 4, pp. 4064–4067.
- [11] F. Yang, M. Paindavoine, H. Abdi, and A. Monopoly, "Development of a fast panoramic face mosaicing and recognition system," *Opt. Eng.*, vol. 44, no. 8, pp. 087 005/1–087 005/10, 2005.
- [12] X. Liu and T. Chen, "Geometry-assisted statistical modeling for face mosaicing," in *Proc. IEEE Int. Conf. Image Process.*, Sep. 2003, vol. 2, pp. 883–886.
- [13] X. Liu and T. Chen, "Pose-robust face recognition using geometry assisted probabilistic modeling," in *Proc. Int. Conf. Comput. Vis. and Pattern Recog.*, Jun. 2005, vol. 1, pp. 502–509.
- [14] T. Sim, S. Baker, and M. Bsat, "The CMU pose, illumination, and expression database," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 12, pp. 1615–1618, Dec. 2003.
- [15] R. Singh, M. Vatsa, A. Ross, and A. Noore, "Performance enhancement of 2D face recognition via mosaicing," in *Proc. 4th IEEE Workshop Autom. Identification Adv. Technol.*, 2005, pp. 63–68.
- [16] S. B. Kang, "A survey of image-based rendering techniques," Cambridge Res. Lab., Cambridge, MA, Tech. Rep. CRL 97/4, 1997.
- [17] K. W. Bowyer, K. Chang, and P. Flynn, "A survey of 3D and multi-modal 3D+2D face recognition," Dept. Comput. Sci. Eng., Notre Dame, IN, Tech. Rep. TR-2004-22-9, 2004.
- [18] T. Serre, L. Wolf, and T. Poggio, "Object recognition with features inspired by visual cortex," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. and Pattern Recog.*, Jun. 2005, vol. 2, pp. 994–1000.
- [19] M. Riesenhuber and T. Poggio, "Hierarchical models of object recognition in cortex," *Nat. Neurosci.*, vol. 2, no. 11, pp. 1019–1025, 1999.
- [20] C. Xu and J. L. Prince, "Snakes, shapes, and gradient vector flow," *IEEE Trans. Image Process.*, vol. 7, no. 3, pp. 359–369, Mar. 1998.
- [21] P. J. Burt and E. H. Adelson, "A multiresolution spline with application to image mosaics," *ACM Trans. Graph.*, vol. 2, no. 4, pp. 217–236, 1983.
- [22] S. Periaswamy and H. Farid, "Elastic registration in the presence of intensity variations," *IEEE Trans. Med. Imag.*, vol. 22, no. 7, pp. 865–874, Jul. 2003.
- [23] D. Hill, C. Studholme, and D. Hawkes, "Voxel similarity measures for automated image registration," in *Proc. 3rd SPIE Conf. Vis. Biomed. Comput.*, 1994, pp. 205–216.
- [24] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality image registration by maximization of mutual information," *IEEE Trans. Med. Imag.*, vol. 16, no. 2, pp. 187–198, Apr. 1997.
- [25] B. Duc, S. Fischer, and J. Bigun, "Face authentication with Gabor information on deformable graphs," *IEEE Trans. Image Process.*, vol. 8, no. 4, pp. 504–516, Apr. 1999.
- [26] C. Liu and H. Wechsler, "Gabor feature based classification using the enhanced fisher linear discriminant model (EFM) for face recognition," *IEEE Trans. Image Process.*, vol. 11, no. 4, pp. 467–476, Apr. 2002.
- [27] D. J. Field, "Relations between the statistics of natural images and the response properties of cortical cells," *J. Opt. Soc. Amer.*, vol. 4, no. 12, pp. 2379–2394, Dec. 1987.
- [28] H. G. Chew, C. C. Lim, and R. E. Bogner, "An implementation of training dual-nu support vector machines," in *Optimization and Control With Applications*, L. Qi, K. Teo, and X. Yang, Eds. Norwell, MA: Kluwer, 2004.
- [29] P. H. Chen, C. J. Lin, and B. Schölkopf, "A tutorial on ν -support vector machines," *Appl. Stoch. Models Bus. Ind.*, vol. 21, no. 2, pp. 111–136, 2005.
- [30] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed. Hoboken, NJ: Wiley, 2000.
- [31] R. Singh, M. Vatsa, and A. Noore, "Hierarchical fusion of multi-spectral face images for improved recognition performance," *Inf. Fusion*, 2008. DOI: 10.1016/j.inffus.2006.002.
- [32] X. Chen, P. Flynn, and K. Bowyer, "IR and visible light face recognition," *J. Comput. Vis. Image Underst.*, vol. 99, no. 3, pp. 332–358, 2005.
- [33] S. Kong, J. Heo, B. Abidi, J. Paik, and M. Abidi, "Recent advances in visual and infrared face recognition—A review," *J. Comput. Vis. Image Underst.*, vol. 97, no. 1, pp. 103–135, 2005.
- [34] F. Prokoshi, "History, current status, and future of infrared identification," in *Proc. IEEE Workshop Comput. Vis. Beyond Visible Spectrum: Methods and Appl.*, 2000, pp. 5–14.
- [35] S. Z. Li, R. F. Chu, M. Ao, L. Zhang, and R. He, "Highly accurate and fast face recognition using near infrared images," in *Proc. IAPR Int. Conf. Biometrics*, Jan. 2006, pp. 151–158.
- [36] M. Turk and A. Pentland, "Eigenfaces for recognition," *J. Cogn. Neurosci.*, vol. 3, no. 1, pp. 71–86, 1991.
- [37] P. N. Belhumeur, J. P. Hespanha, and D. J. Kriegman, "Eigenfaces vs. Fisherfaces: Recognition using class specific linear projection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 19, no. 7, pp. 711–720, Jul. 1997.
- [38] P. Penev and J. Atick, "Local feature analysis: A general statistical theory for object representation," *Netw.: Comput. Neural Syst.*, vol. 7, no. 3, pp. 477–500, 1996.
- [39] R. Singh, M. Vatsa, and A. Noore, "Textural feature based face recognition for single training images," *Electron. Lett.*, vol. 41, no. 11, pp. 640–641, May 2005.



Richa Singh (S'04) received the M.S. degree in computer science in 2005 and is currently working toward the Ph.D degree in computer science at West Virginia University, Morgantown.

From July 2002 to July 2004, she was actively involved in the development of a multimodal biometric system, which includes face, fingerprint, and iris recognition at the Indian Institute of Technology Kanpur. She is a Graduate Research Assistant with the Lane Department of Computer Science and Electrical Engineering, West Virginia University. She has

more than 55 publications in refereed journals, book chapters, and conference proceedings. Her current research interests are pattern recognition, image processing, machine learning, granular computing, biometric authentication, and data fusion.

Ms. Singh is a member of the Phi Kappa Phi, Eta Kappa Nu, and Upsilon Pi Epsilon honor societies. She was the recipient of four best paper awards.



Mayank Vatsa (S'04) received the M.S. degree in computer science in 2005 and is currently working toward the Ph.D degree in computer science at West Virginia University, Morgantown.

From July 2002 to July 2004, he was actively involved in the development of a multimodal biometric system, which includes face, fingerprint, and iris recognition at the Indian Institute of Technology Kanpur. He is a Graduate Research Assistant with the Lane Department of Computer Science and Electrical Engineering, West Virginia University. He has

more than 55 publications in refereed journals, book chapters, and conference proceedings. His current research interests are pattern recognition, image processing, uncertainty principles, biometric authentication, watermarking, and information fusion.

Mr. Vatsa is a member of the Phi Kappa Phi, Sigma Xi, Eta Kappa Nu, and Upsilon Pi Epsilon honor societies. He was the recipient of four best paper awards.



Arun Ross (M'04) received the B.E. (Hons.) degree in computer science from the Birla Institute of Technology and Science, Pilani, India, in 1996 and the M.S. and Ph.D. degrees in computer science and engineering from Michigan State University, East Lansing, in 1999 and 2003, respectively.

Between 1996 and 1997, he was with the Design and Development Group, Tata Elxsi (India) Ltd., Bangalore, India. He also spent three summers (2000–2002) with the Imaging and Visualization Group, Siemens Corporate Research, Inc., Princeton,

NJ, working on fingerprint recognition algorithms. He is currently an Assistant Professor with the Lane Department of Computer Science and Electrical Engineering, West Virginia University, Morgantown, where he is actively involved in the development of pattern recognition and biometrics curricula. He was an invited speaker at the Kavli Frontiers of Science Symposium that was organized by the National Academy of Sciences in November 2006. He is a coauthor of the book *Handbook of Multibiometrics* (Springer, 2006). His research interests include pattern recognition, classifier combination, machine learning, computer vision, and biometrics.



Afzel Noore (M'03) received the Ph.D. degree in electrical engineering from West Virginia University, Morgantown.

He was a Digital Design Engineer with Philips India. From 1996 to 2003, he served as the Associate Dean for Academic Affairs and Special Assistant to the Dean in the College of Engineering and Mineral Resources, West Virginia University. He is currently a Professor with the Lane Department of Computer Science and Electrical Engineering, West Virginia University. His research interests include computa-

tional intelligence, biometrics, software reliability modeling, machine learning, hardware description languages, and quantum computing. His research has been funded by NASA, NSF, Westinghouse, GE, Electric Power Research Institute, the U.S. Department of Energy, and the U.S. Department of Justice. He has more than 80 publications in refereed journals, book chapters, and conference proceedings.

Dr. Noore is a member of the Phi Kappa Phi, Sigma Xi, Eta Kappa Nu, and Tau Beta Pi honor societies. He was the recipient of four best paper awards.