Optimization of Color Conversion for Face Recognition

Creed F. Jones III

Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University,
Blacksburg, VA 24061-0111, USA

Department of Computer Science, Seattle Pacific University, Seattle, WA 98119-1957, USA
Email: crjones4@vt.edu

A. Lynn Abbott

Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University,
Blacksburg, VA 24061-0111, USA
Email: abbott@vt.edu

This paper concerns the conversion of color images to monochromatic form for the purpose of human face recognition. Many face recognition systems operate using monochromatic information alone even when color images are available. In such cases, simple color transformations are commonly used that are not optimal for the face recognition task. We present a framework for selecting the transformation from face imagery using one of three methods: Karhunen-Loève analysis, linear regression of color distribution, and a genetic algorithm. Experimental results are presented for both the well-known eigenface method and for extraction of Gabor-based face features to demonstrate the potential for improved overall system performance. Using a database of 280 images, our experiments using these methods resulted in performance improvements of approximately 4% to 14%.

Keywords and phrases: face recognition, color image analysis, color conversion, Karhunen-Loève analysis.

1. INTRODUCTION

Most single-view face recognition systems operate using intensity (monochromatic) information alone. This is true even for systems that accept color imagery as input. The reason for this is not that multispectral data is lacking in information content, but often because of practical considerations—difficulties associated with illumination and color balancing, for example, as well as compatibility with legacy systems. Associated with this is a lack of color image databases with which to develop and test new algorithms. Although work is in progress that will eventually aid in color-based tasks (e.g., through color constancy [1]), those efforts are still in the research stage.

When color information is present, most of today's face recognition systems convert the image to monochromatic form using simple transformations. For example, a common mapping [2, 3] produces an intensity value I_i by taking the average of red, green, and blue (RGB) values (I_r, I_g, and I_b, resp.):

$$I_i(x, y) = \frac{I_r(x, y) + I_g(x, y) + I_b(x, y)}{3}. \qquad (1)$$

The resulting image is then used for feature extraction and analysis.

We argue that more effective system performance is possible if a color transformation is chosen that better matches the task at hand. For example, the mapping in (1) implicitly assumes a uniform distribution of color values over the entire color space. For a task such as face recognition, color values tend to be more tightly confined to a small portion of the color space, and it is possible to exploit this narrow concentration during color conversion. If the transformation is selected based on the expected color distribution, then it is reasonable to expect improved recognition accuracies.

This paper presents a task-oriented approach for selecting the color-to-grayscale image transformation. Our intended application is face recognition, although the framework that we present is applicable to other problem domains.

We assume that frontal color views of the human face are available, and we develop a method for selecting alternate weightings of the separate color values in computing a single monochromatic value. Given the rich color content of the human face, it is desirable to maximize the use of this content even when full-color computation and matching is not used. As an illustration of this framework, we have used the Karhunen-Loève (KL) transformation (also known as principal components analysis) of observed distributions in the color space to determine the improved mapping.

Other work [4] has suggested that alternative color spaces provide no real benefit for locating skin in images because these spaces do not increase the separability of the skin and nonskin classes. However, to extract features for face recognition, we do not wish to discriminate skin from nonskin regions, but rather to extract meaningful image features within the skin area. Queisser [5] used the properties of color distributions of a set of similar images to select a new color space for object classification. Abbott and Zhao [6, 7] developed a color-space quantization approach for the recognition of naturally textured objects, but did not consider that for face recognition. Torres has demonstrated that color information can provide additional accuracy for the "eigenface" approach [8], although there is no discussion of optimal color representation. Heseltine et al. [9] measured the performance effect on eigenface-based face recognition of a number of preprocessing techniques, including several color transformations (RGB to hue, brightness-insensitive hue, etc.) and found that these color methods actually degraded the recognition accuracy. However, the techniques that they explored were general color transformations that were not based on the content of the images.

The remainder of this paper is organized as follows. Section 2 presents our approach for using KL analysis to determine a suitable single color axis for a given set of RGB images, and Section 3 presents experimentally derived color transformation data using this method. In Section 4, we investigate the use of KL analysis on color data in CIE L-a-b format. Section 5 describes an alternative method based on linear regression analysis of RGB pixel data, while Section 6 discusses our experimental use of a genetic algorithm to select the color conversion. Section 7 presents the face recognition accuracy improvement observed with the eigenface method from using the KL derived color transformation, and Section 8 describes the effect of the optimal color conversion on feature vectors ("jets") extracted using complex Gabor filters. Finally, Section 9 presents concluding remarks.

2. KL COLOR CONVERSION—RGB

Pixels in the original color image can be represented as the vector $I(x, y) = [I_r(x, y) \ I_g(x, y) \ I_b(x, y)]^T$, where the r, g, and b subscripts denote the red, green, and blue color planes, respectively. As described in (1), face recognition systems typically use an intensity plane derived as $I_i(x, y) = 1/3[1 \ 1 \ 1]I(x, y)$. We propose that human face images exhibit common characteristics that can be exploited in the conversion from a full-color representation to a monochrome image. In the hue-saturation plane, for example, face pixels from a mixture of ethnic groups are well clustered [10], with only the intensity plane varying markedly. This suggests that the standard intensity plane is in fact more sensitive to variation due to ethnic type, which is undesirable.

To determine an improved linear transformation, we want to find the optimum transformation vector w such that $M(x, y) = w^T I(x, y)$, where I is the original color image and M is the resulting single-plane image. We make the assumption that the optimum transformation corresponds closely to the expected distribution of pixel values within the original color space. With this in mind, it is possible to select w by using the KL transformation to determine the projection with uncorrelated axes. The resulting color space has been called the "Karhunen-Loève color space" for an unspecified pixel population [11, 12]; here, we specifically restrict it to the face area. For a given distribution of pixel values, the eigenvector corresponding to the largest eigenvalue defines the direction along which the data is the least correlated, and therefore most likely to be of use in recognition tasks.

The KL transformation is determined from the covariance matrix of the distribution. For this application, the input datum is the ensemble of pixel values from a set of training images, taken from the region containing the face. We form the covariance matrix S as follows:

$$
S = \frac{1}{N} \begin{bmatrix} \sum_m p_r{}^2 & \sum_m p_r p_g & \sum_m p_r p_b \\ \sum_m p_g p_r & \sum_m p_g{}^2 & \sum_m p_g p_b \\ \sum_m p_b p_r & \sum_m p_b p_r & \sum_m p_b{}^2 \end{bmatrix}
$$
$$
- \frac{1}{N^2} \begin{bmatrix} \left(\sum_m p_r\right) \\ \left(\sum_m p_g\right) \\ \left(\sum_m p_b\right) \end{bmatrix} \begin{bmatrix} \left(\sum_m p_r\right) \\ \left(\sum_m p_g\right) \\ \left(\sum_m p_b\right) \end{bmatrix}^T ,
$$

(2)

where p is the collection of N color pixel vectors. The KL transformation is then given by the eigenvectors $\{u_i\}$ of S, concatenated into the matrix $U = [u_1 \ u_2 \ u_3]$. The eigenvector u_1, associated with the largest eigenvalue, is of primary interest here; it represents the direction of most variability in the data within the original space. Projection of RGB values onto this axis represents a color-to-grayscale conversion with the highest potential for discrimination.

The normalization of the conversion vector w requires consideration. A unit vector will, by definition, not change the magnitude of the vector quantity that it operates on. However, this is not appropriate for conversion of three-component color quantities (where each component can range up to full scale) to monochrome, since any three-color vector with magnitude greater than unity will saturate in the monochrome plane. We prevent saturation by normalizing the vector having RGB components at full scale to a magnitude of 1. Therefore, the conversion vectors that we compute are normalized by $\sqrt{3}$.

3. RESULTS OF KL ANALYSIS ON RGB DATA

The images used in this study are frontal-view, color face images from two databases (described in [13, 14]). Each image is of size 240 rows by 300 columns. Prior to this study, the images were spatially registered so that the centers of the eye sockets are at fixed locations, the line between the eye centers is horizontal, and the distance between

eye centers is 60 pixels, in accordance with developing standards for face recognition image interchange [15]. No effort was made to color-correct or contrast-equalize the images. To determine the color conversion that is most suited for the face features, we process only a portion of the face image that represents the area of the face with minimal included background and hair. The extent to be processed, a region 90 pixels wide by 140 pixels high, is indicated in Figure 1.

The KL analysis described in Section 2 yields an eigenvector u_1 describing the axis of projection with the largest variance in the original data, which we call the conversion vector. Let the three components of this vector be represented by $u_1 = [u_{11} \quad u_{12} \quad u_{13}]^T$. Because this vector has unit length $r = \sqrt{u_{11}^2 + u_{12}^2 + u_{13}^2} = 1$, we can represent it using spherical coordinates and completely describe the color mapping by the two angular quantities θ and ϕ:

$$\phi = \arccos(u_{13}), \qquad \theta = \arccos\left(\frac{u_{11}}{\sin(\phi)}\right). \qquad (3)$$

To illustrate the meaningfulness of the transformation, several scatter diagrams are shown in Figure 2. Four collections of face images are represented as well as some natural images of random content. For each image, the color histogram was computed and the conversion vector u_1 obtained. The resulting conversion vectors are indicated as points in $[\phi, \theta]$ space. Each face image collection consists of several sets of 21 images, each for a single individual. The natural images contain a mix of object types including landscapes, photographs of sporting events, and astronomy images.

It can be seen that the optimal color conversion vectors u_1 computed for the face images are distinct from those for more general natural images, indicating that the red, green, and blue color planes carry different degrees of information for the specific class of face images. The figure also indicates the position in this space of an equal-weighted color conversion, which appears to represent a good estimate for the optimal conversion for general natural images, but is not well suited for the face image collections. The selection of face databases used in our testing contain color distributions that generally correspond to $[\phi = 1.01, \theta = 0.662]$, which in turn corresponds to a conversion vector of

$$u_1 = \frac{1}{\sqrt{3}} \begin{bmatrix} \sin\phi\cos\theta \\ \sin\phi\sin\theta \\ \cos\phi \end{bmatrix} = \begin{bmatrix} 0.3858 \\ 0.3004 \\ 0.3070 \end{bmatrix}. \qquad (4)$$

This should be compared with the equal-weighted values of $[0.333 \quad 0.333 \quad 0.333]^T$. We observe that the KL procedure for these images results in a color space that more heavily weights the red color component than the green and blue. This indicates that face images contain more uncorrelated variation in the red plane than in the green or blue planes.

Note that the preceding eigenvalue-eigenvector analysis concerns only the color-to-monochrome conversion process, and is independent of the face recognition approach that is
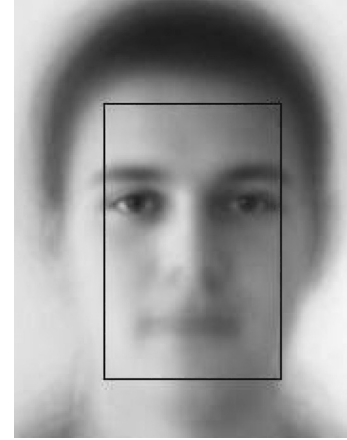

FIGURE 1: Illustration of image extent to be processed for color conversion and recognition. This monochrome image is an "average" image.

KL face DB 1 ・   KL natural images ▪
KL face DB 2 ▲   Equal-weighted RGB ●
KL face DB 3 ✕   Line-fit face DB 1 +
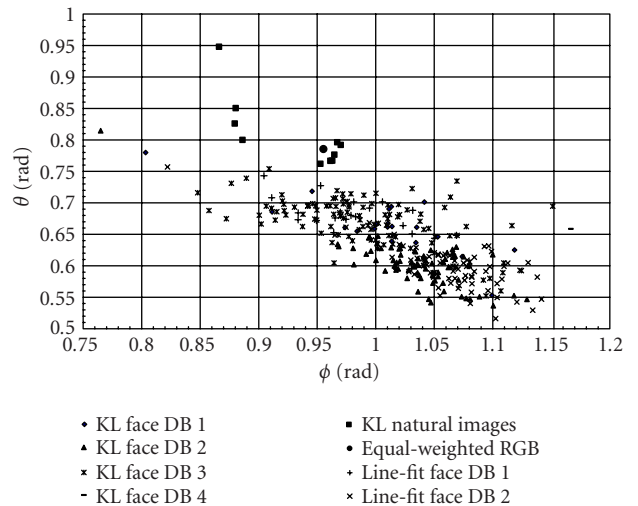KL face DB 4 −   Line-fit face DB 2 ✕

FIGURE 2: Principal component directions, using spherical coordinates $[\phi, \theta]$, for several histograms in RGB space. Each point in the diagram represents the orientation of greatest variability in the aggregate color histogram of a complete image database. The "line-fit" cases listed in the legend are described in Section 5.

used. We propose that *any* face recognition technique could benefit from a careful examination of the initial conversion from color to monochrome images.

4. KL COLOR CONVERSION—L-a-b

RGB is not always the most convenient space in which to process color information. The CIE tristimulus system represents a color in terms of its three coordinates relative to a reference color, usually a standard illuminant [16]. However, equal distances in the XYZ space are perceived as unequal, so the L-a-b color space is defined so that color distances are perceived as linear.

The L-a-b space is defined as follows [16]:

$$p_{\mathrm{L}} = 116\left(\frac{p_Y}{Y_0}\right)^{1/3} - 16,$$

$$p_{\mathrm{a}} = 500\left[\left(\frac{p_X}{X_0}\right)^{1/3} - \left(\frac{p_Y}{Y_0}\right)^{1/3}\right], \qquad (5)$$

$$p_{\mathrm{b}} = 200\left[\left(\frac{p_Y}{Y_0}\right)^{1/3} - \left(\frac{p_Z}{Z_0}\right)^{1/3}\right],$$

where

$$\begin{bmatrix} I_X \\ I_Y \\ I_Z \end{bmatrix} = \begin{bmatrix} 0.412453 & 0.357580 & 0.189423 \\ 0.212671 & 0.715160 & 0.072169 \\ 0.019334 & 0.119193 & 0.950227 \end{bmatrix} \begin{bmatrix} I_r \\ I_g \\ I_b \end{bmatrix} \qquad (6)$$

for the D65 standard illuminant used as the color reference point ($[X_0 \; Y_0 \; Z_0] = [1004.26 \; 1056.79 \; 1150.71]$).

Our KL-based approach for selecting the color conversion produces a linear transformation of the RGB color values; thus, we could expect that using the KL process on the XYZ values would produce the same result within computation accuracy. However, the relation between RGB and L-a-b is nonlinear, and the L-a-b space is in some sense more relevant to human perception, so that application of the KL procedure defined in Section 2 would be expected to produce useful results.

In fact, as can be seen in Figure 3, the KL transformation on L-a-b data does not yield distinctive data for face pixels as opposed to image pixels from more general scenes. This suggested that the "optimal" color conversion obtained from L-a-b data does not provide any beneficial added feature content. Experimentation with the eigenvalues of face images converted to L-a-b representation, and then projected onto the axis found by using KL on the resulting histogram data (as described in Section 7) showed that this was the case; information contained in the most significant n axes was not greater (and in fact frequently less) than that for the L plane of the corresponding L-a-b images. It is possible that a transformation resulting in a linear perception of color distance inherently concentrates useful detail information in the L plane.

5. COLOR CONVERSION THROUGH LINEAR REGRESSION

Queisser discusses (in [5]) the use of a least-squared-error line-fit to RGB data to define a new color axis that is best suited to images of a particular class of object. In his study, images of wood panels and food products were shown to be more suited for object detection and inspection in the resulting single-color plane than in any of the hue-saturation-intensity (HSI) axes. The other axes relate to additional magnitude and chromaticity information.

We consider a similar approach in the RGB space. We performed least-squared-error fits to our RGB data with the added constraint that the new axis of projection should pass through the RGB origin. The purpose of this is to force a pixel
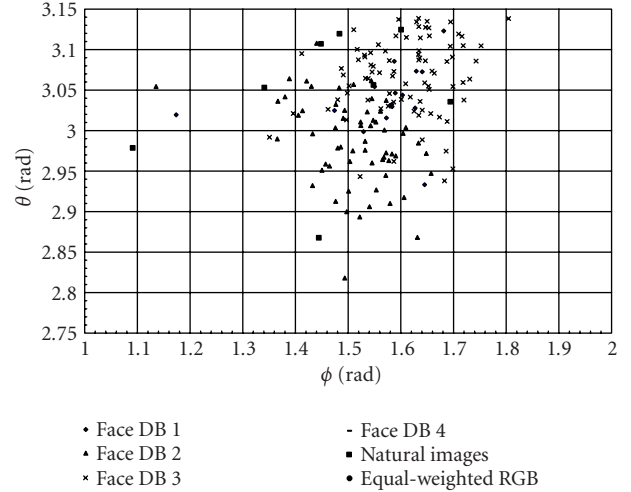

FIGURE 3: Principal component directions, using spherical coordinates $[\phi, \theta]$, for several histograms in L-a-b space. Each point in the diagram represents the orientation of greatest variability in the aggregate color histogram of a complete image database. The "line-fit" cases listed in the legend are described in Section 5.

with zero in all color planes to map to a black pixel in the new space. The transformation matrix is as follows:

$$\begin{bmatrix} \beta \\ s \\ t \end{bmatrix} = \begin{bmatrix} \bar{r} & \bar{g} & \bar{b} \\ \dfrac{-\bar{g}}{\sqrt{\bar{r}^2 + \bar{g}^2}} & \dfrac{\bar{r}}{\sqrt{\bar{r}^2 + \bar{g}^2}} & 0 \\ \dfrac{-\bar{r}\bar{b}}{\sqrt{\bar{r}^2 + \bar{g}^2}} & \dfrac{-\bar{g}\bar{b}}{\sqrt{\bar{r}^2 + \bar{g}^2}} & \dfrac{\bar{r}^2 + \bar{g}^2}{\sqrt{\bar{r}^2 + \bar{g}^2}} \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}. \qquad (7)$$

Applying (7) to the pixel data from the face box areas of the sample databases, we obtain the data presented in Figure 2 as the "line-fit" data. As before, we are only interested in the primary axis, β in this transformation. The results are very similar to those obtained by the KL method, but with much lower computational cost because only the red, green, and blue sample means are required.

6. COLOR CONVERSION THROUGH GENETIC ALGORITHM SEARCH

To further investigate the determination of the color projection by optimizing the face recognition accuracy, we applied a genetic algorithm to the color vector selection process. Each individual in the population consisted of a $[\phi, \theta]$ pair as defined in (3). The optimization algorithm had the following properties:

  (i) population size of 100;
  (ii) breeding by averaging of $[\phi, \theta]$ values;
  (iii) population initialized with random values;
  (iv) "roulette wheel" selection model with elitism (the best two candidates in each generation will persist [17]);

Table 1: Results of genetic algorithm.

Iteration	1	2	$\cdots$	99	100
Best ϕ	1.5069	1.3899	$\cdots$	1.1072	1.1072
Best θ	0.7341	0.4707	$\cdots$	0.6496	0.6496
Error	0.0965	0.0961	$\cdots$	0.0948	0.0948

(v) mutation by perturbation of a random individual (probability of mutation was 0.005);

(vi) error function to be minimized was the difference between 1.0 and the sum of the first 8 normalized eigenvalues.

This study therefore attempted to maximize the performance of the face recognition system as simulated by the sum of the largest 8 eigenvalues.

The results of 100 generations of testing on one sample database are summarized in Table 1. The testing data suggested that the error surface was slowly changing and not unimodal. After only 100 iterations, the convergence is clearly dominated by the effect of mutation rather than breeding. The resulting vector differs in error from the results obtained by KL computation by only 0.00077.

The advantage of the genetic algorithm for this purpose is its flexibility in that it is possible to define the error function in terms of any computable metric of overall system performance. For example, this could be biased toward a particular combination of Type I (false positive) and Type II (false negative) errors of recognition on a given database. The major disadvantage of this method is its computational requirements. For a relatively modest database of fifty individuals, 100 generations took more than six hours to run on a 1 GHz Pentium III machine. In addition, the genetic algorithm has unpredictable convergence behavior and a set of performance parameters that may require tuning. Our experimentation with a GA roughly confirmed the earlier computed results.

7. EFFECT OF OPTIMIZED COLOR CONVERSION ON FACE RECOGNITION ACCURACY

To evaluate the effect of our color conversion method on face recognition accuracy, we considered the effect on performance of the well-known eigenface method [18, 19]. This technique uses principal components analysis of a collection of face images, treated as one-dimensional vectors, to determine the linear combinations of pixel locations that form the best projective axes for the collection. Early work in this area focused on the use of a small set of these projections to adequately represent a face image, while later work (beginning around 1990) applied this same technique to recognition. The new "face space" defined by the most significant basis vectors, called "eigenfaces," is used for pattern recognition based on a distance measure.

For any principal component analysis, the ratio of an eigenvalue to the sum of all the eigenvalues is proportional to the mean squared error implied by exclusion of the cor-

responding eigenvector [20]. Thus, we can examine the cumulative sum of eigenvalues 1 through n, plotted versus n, to compare the information contained in the first n eigenfaces (the "principal components"). In this way, we can predict the performance of the eigenface method on the two databases. Table 2 shows the individual and cumulative eigenvalues for a typical database of face images.

Figure 4 shows a plot of the cumulative eigenvalues, which gives a measure of the accuracy achievable by truncating all higher eigenvalues. Using the optimized color conversion produces a modest, yet consistent, improvement in the potential accuracy. The increased information is more pronounced for the more significant eigenvalues.

By comparison, we also evaluated the magnitude of the initial eigenvectors for the eigenface method when using the line-fit method described in Section 5. The cumulative eigenvalues computed by using the β axis as the new image plane are shown in Table 3 and exhibit a similar increase in information in the lowest eigenfaces. In fact, for all of the databases we examined, the use of the line fit gave essentially equal performance as measured by the normalized eigenvalues.

For confirmation of these predictions of increased performance, we measured the face recognition accuracy on a complete eigenface recognition implementation. We will not describe the specifics of the eigenface method here as they are covered well in [18, 19]. For our test, a training phase and a test phase were implemented. The training phase computes the desired transformation by solving for the eigenvalues of the matrix composed of the concatenation of the training images. Testing is performed by applying this transformation to a set of probe images of the same individuals and measuring the Euclidean distance from the probe image data to the exemplars of each individual, defined as the average in "face space" of each training image of that individual. The probe images were not present in the training set. Note that the eigenface implementation was fairly simplistic; our objective was not to achieve overall high recognition accuracy but to measure the effect of using our color conversion.

To measure the performance in a consistent fashion, we adopted the method used in the NIST FERET studies [21]. The results for each probe image are ranked in order of increasing Euclidean distance. The performance score for a particular experiment R_n is defined to be the ratio of the number of times that the correct identity is in the top n candidates (the n nearest exemplars to the probe image) to the total number of probe images tested.

Table 4 summarizes the eigenface performance for three values of n (2, 5, and 10) for a particularly difficult database of 280 images. Many of the images exhibit poor contrast, and there is significant variation in expression by the human subjects. Two sets of results are shown: the first for a typical equal-weighted conversion from RGB to monochrome and the second for a transformation vector derived using the KL procedure described above. The results show significant improvements in performance scores (roughly in the range of 6% to 14%) when the KL conversion was used. Although the database was relatively small, and therefore care must be

TABLE 2: Eigenvalues for a typical face database using the KL method to determine the RGB conversion.

	Index$_i$	1	2	3	4	5	6	7	8
Equal-weighted RGB conversion	Eigenvalue λ_i	0.51613	0.13616	0.06179	0.05184	0.04326	0.03875	0.02879	0.02683
	Cumulative $\sum \lambda_i$	0.51613	0.65229	0.71408	0.76592	0.80918	0.84793	0.87672	0.90354
KL-computed RGB conversion	Eigenvalue λ_i	0.53441	0.13210	0.05929	0.05223	0.04181	0.03812	0.02850	0.02488
	Cumulative $\sum \lambda_i$	0.53441	0.66651	0.72580	0.77803	0.81984	0.85796	0.88646	0.91134

TABLE 3: Eigenvalues for a typical face database using the line-fit method to determine the RGB conversion.

	Index$_i$	1	2	3	4	5	6	7	8
Equal-weighted RGB conversion	Eigenvalue λ_i	0.51613	0.13616	0.06179	0.05184	0.04326	0.03875	0.02879	0.02683
	Cumulative $\sum \lambda_i$	0.51613	0.65229	0.71408	0.76592	0.80918	0.84793	0.87672	0.90354
Line-fit RGB conversion	Eigenvalue λ_i	0.53280	0.13175	0.05801	0.05210	0.04064	0.03644	0.02898	0.02512
	Cumulative $\sum \lambda_i$	0.53280	0.66445	0.72256	0.77467	0.81531	0.85174	0.88073	0.90585
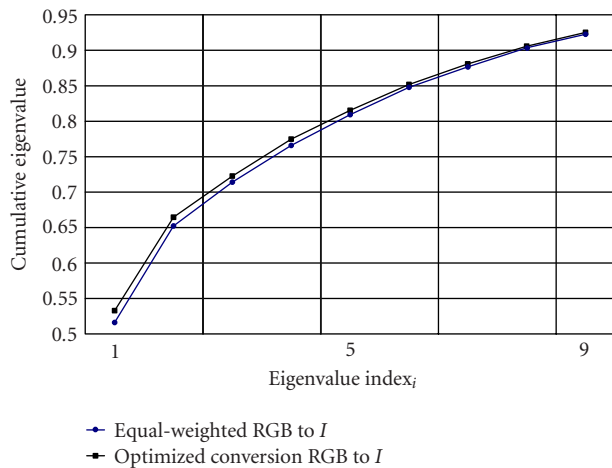
FIGURE 4: Comparison of cumulative eigenvalues for the eigenface procedure. The optimized RGB to monochrome conversion results in more significant information in the first n eigenfaces.

taken in extrapolating these accuracy values to larger sets, they provide a strong indication that the color conversion process can have a sizable impact on face recognition performance.

Because the face images had a noticeable increase in contrast as a result of the KL derived RGB to monochrome transformation, there was a concern that the KL derived method was doing no more than could be obtained from a common histogram equalization on the color image. To explore this idea, the eigenface performance was also measured with and without the use of a histogram equalization preprocessing step. Each color plane in the original RGB space was enhanced using a standard 256-to-64-bin histogram flattening procedure. The results show that, rather than a performance increase similar to that obtained from the optimized color conversion, the histogram equalization actually produced a

severe decrease in accuracy. It is believed that this is due to the global nature of the process, which may have resulted in a suppression of the facial features that are useful for recognition. We conclude that color histogram equalization is not a useful preprocessing step for eigenface face recognition, regardless of the choice of method for color transformation.

8. EFFECT ON FACE FEATURE DISCRIMINABILITY USING GABOR FILTERS

Another technique for face recognition is based on the application of a family of Gabor filters to monochrome face images [22, 23, 24]. A two-dimensional Gabor filter is a directed complex sinusoid in the image plane, decaying exponentially as a function of distance from the filter's origin. At a set of preselected locations on the face, Gabor filters at various related directions and sinusoidal frequencies are applied and the complex responses are assembled into a feature vector known as a "jet." Several techniques exist for performing face recognition using these jets.

We have evaluated the potential improvement in Gabor-based methods from the use of an optimized color transformation by evaluating the relative distances between the Gabor jets from the same point on different faces, with and without the use of the KL derived color transformation. To obtain the interjet distances, we consider the jets as 80-element vectors and determine the Mahalanobis distance by the usual method. To measure the effectiveness of a set of jets for face discrimination, we consider (at each facial landmark) the ratio of the minimum interjet distance between two different faces to the maximum interjet distance, as well as the ratio of the minimum interjet distance to the average of all interjet distances for that landmark. Ratios were used to provide some normalization.

When the KL derived color transformation is used, the min-to-max ratio improved by 4.4% on a set of ten facial landmarks over the test database, while the min-to-average

TABLE 4: Improvement in face recognition performance with new color conversion procedure. The monochrome eigenface recognition procedure was used on a database of 280 color images. The second and third columns show the recognition accuracy values that were obtained when the images were color-converted with the standard equal-weight method and with our KL method, respectively.

	Equal-weighted RGB	KL RGB	Improvement
R_2—probe in top 2	0.343	0.371	8.3%
R_5—probe in top 5	0.514	0.600	13.8%
R_{10}—probe in top 10	0.686	0.728	6.2%

ratio increased by 6%. Interestingly, the average interjet distance actually decreased slightly, indicating that the minimum interjet distances were larger than when the usual monochrome intensity images were used. This is an initial indication that Gabor-based methods may have greater discrimination between different individuals when the KL derived color-to-monochrome transformation is used, since the underlying features are more distinctive.

9. CONCLUSIONS

This paper has presented a new approach for converting color images to monochromatic form. By tailoring the conversion process to the needs of a particular task, such as human face recognition, it is possible to improve the overall system performance.

Most existing face recognition systems operate using monochromatic information alone, even when color information is available. In such cases, a simple and suboptimal conversion process is typically used. We argue that recognition accuracies can be improved if the color-conversion process is selected based on the expected color distributions. We explored three such approaches to determine an improved mapping empirically: Karhunen-Loève analysis of the color pixel distributions, a least-squared-error line fit in RGB space, and a genetic algorithm.

The color-conversion method presented in this paper is independent of the actual face recognition approach that is used. For testing purposes, however, we have used the well-known eigenface method. Our experiments using the eigenface method for recognition resulted in performance improvements in the range of approximately 6% to 14% for a database of 280 color images. Relative distance measurements of Gabor jets of the face area also showed an increase in discriminability of 4% to 6%. Evaluation of the cumulative eigenvalues produced by an eigenface analysis of intensity images and images converted to grayscale form using the computed conversion vector showed a modest yet consistent improvement in the potential accuracy in retaining only the most important n basis vectors.

REFERENCES

[1] G. D. Finlayson, S. D. Hordley, and P. M. Hubel, "Color by correlation: a simple, unifying framework for color constancy," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 23, no. 11, pp. 1209–1221, 2001.

[2] R. Gonzalez and R. Woods, *Digital Image Processing*, Addison-Wesley, NY, USA, 1st edition, 1992.

[3] R. Hunt, "Why is black and white so important in colour?," in *Colour Imaging: Vision and Technology*, L. MacDonald and M. Luo, Eds., John Wiley & Sons, NY, USA, 1999.

[4] A. Albiol, L. Torres, and E. Delp, "Optimum color spaces for skin detection," in *Proc. IEEE International Conference on Image Processing (ICIP '01)*, vol. 1, pp. 122–124, Thessaloniki, Greece, October 2001.

[5] A. Queisser, "Color spaces for inspection of natural objects," in *Proc. IEEE International Conference on Image Processing (ICIP '97)*, vol. 3, pp. 42–45, Washington, DC, USA, October 1997.

[6] A. L. Abbott and Y. Zhao, "Adaptive quantization of color space for recognition of finished wooden components," in *Proc. 3rd IEEE Workshop on Applications of Computer Vision (WACV '96)*, pp. 252–257, Sarasota, Fla, USA, December 1996.

[7] Y. Zhao, "A color identification system based on class-oriented adaptive color space quantization," M.S. thesis, Bradley Department of Electrical and Computer Engineering, Virginia Tech, Blacksburg, Va, USA, 1996.

[8] L. Torres, J. Reutter, and L. Lorente, "The importance of the color information in face recognition," in *Proc. IEEE International Conference on Image Processing (ICIP '99)*, vol. 3, pp. 627–631, Kobe, Japan, October 1999.

[9] T. Heseltine, N. Pears, and J. Austin, "Evaluation of image pre-processing techniques for eigenface-based face recognition," in *Proc. 2nd International Conference on Image and Graphics (ICIG '02)*, Wei Sui, Ed., vol. 4875 of *Proc. SPIE*, pp. 677–685, Hefei, Anhui, China, July 2002.

[10] S. Gong, S. McKenna, and A. Psarrou, *Dynamic Vision: From Images to Face Recognition*, Imperial College Press, London, UK, 2000.

[11] C. Lee, J. Kim, and K. Park, "Automatic human face location in a complex background using motion and color information," *Pattern Recognition*, vol. 29, no. 11, pp. 1877–1889, 1996.

[12] W. Pratt, *Digital Image Processing*, John Wiley & Sons, NY, USA, 1978.

[13] T. Sim, S. Baker, and M. Bsat, "The CMU pose, illumination, and expression (PIE) database of human faces," Tech. Rep. CMU-RI-TR-01-02, Robotics Institute, Carnegie Mellon University, 2001.

[14] L. Spacek, "University of Essex Face Database," 2002, http://cswww.essex.ac.uk/.

[15] P. Griffin, "Face recognition format for data interchange," INCITS M1 Biometrics Standards Committee, Document M1/02-0228, October 2002, http://www.ncits.org/tc_home/m1.htm.

[16] D. MacAdam, *Color Measurement: Theme and Variations*, Springer-Verlag, Berlin, Germany, 1985.

[17] C. De Stefano and A. Marcelli, "Generalization vs. specialization: quantitative evaluation criteria for genetics-based learning systems," in *Proc. International Conference on Systems, Man and Cybernetics (SMC '97)*, vol. 3, pp. 2865–2870, Orlando, Fla, USA, October 1997.

[18] M. Turk and A. Pentland, "Eigenfaces for recognition," *Journal of Cognitive Neuroscience*, vol. 3, no. 1, pp. 71–86, 1991.

[19] P. Grother, "Software Tools for an Eigenface Implementation," 2001, http://www.nist.gov/humanid/feret/.

[20] M. Nadler and E. Smith, *Pattern Recognition Engineering*, John Wiley & Sons, NY, USA, 1993.

[21] P. J. Phillips, H. Moon, S. Rizvi, and P. J. Rauss, "The FERET evaluation methodology for face-recognition algorithms," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 22, no. 10, pp. 1090–1104, 2000.

[22] I. Fasel, M. Bartlett, and J. Movellan, "A comparison of Gabor filter methods for automatic detection of facial landmarks," in *Proc. 5th IEEE International Conference on Automatic Face and Gesture Recognition (FG '02)*, pp. 242–247, Washington, DC, USA, May 2002.

[23] L. Wiskott, J.-M. Fellous, N. Krüger, and C. von der Malsburg, "Face recognition by elastic bunch graph matching," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 19, no. 7, pp. 775–779, 1997.

[24] J. G. Daugman, "Complete discrete 2-D Gabor transforms by neural networks for image analysis and compression," *IEEE Trans. Acoustics, Speech, and Signal Processing*, vol. 36, no. 7, pp. 1169–1179, 1988.

Creed F. Jones III is a faculty member at Seattle Pacific University in Seattle, Washington, where he is an Associate Professor in the Computer Science Department. He received the B.S. and M.S. degrees in electrical engineering from Oakland University, Rochester, Michigan, in 1980 and 1982, respectively, and is a candidate for the Ph.D. degree in computer engineering from Virginia Tech in 2004. From 1982 to 2000, he was an Engineer and Engineering Director with several organizations in the machine vision industry. Mr. Jones' primary research interests involve biometric identification including face recognition, computer vision, and image processing. Mr. Jones is the Chair of the International Committee for Information Technology Standards (INCITS) M1.3, task group for standardization of biometric data formats, and is a member of the IEEE Computer Society.

A. Lynn Abbott is a faculty member at Virginia Tech, Blacksburg, Virginia, where he is an Associate Professor in the Bradley Department of Electrical and Computer Engineering. He received the B.S. degree from Rutgers University in 1980, the M.S. degree from Stanford University in 1981, and the Ph.D. degree from the University of Illinois in 1990, all in electrical engineering. From 1980 to 1985, he was a member of Technical Staff at AT&T Bell Laboratories where his duties involved hardware and software design of data communications equipment. Dr. Abbott's primary research interests involve computer vision and image processing, with emphasis on range estimation and manufacturing automation. He is also interested in pattern recognition, artificial intelligence, and high-performance computer architectures for image processing. Dr. Abbott is a member of the IEEE Computer Society, ACM, Sigma Xi, and the Pattern Recognition Society. He also serves as an Associate Editor for the journal of Computers and Electronics in Agriculture.