Face Recognition under Varying Illumination

ABSTRACT

This paper proposes a novel pipeline to develop a Face Recognition System robust to illumination variation. We consider the case when only one single image per person is available during the training phase. In order to utilize the superiority of Linear Discriminant Analysis (LDA) over Principal Component Analysis (PCA) in regard to variable illumination, a number of new images illuminated from different directions are synthesized from a single image by means of the Quotient Image. Furthermore, during the testing phase, an iterative algorithm is used for the restoration of frontal illumination of a face illuminated from any arbitrary angle. Experimental results on the YaleB database show that our approach can achieve a top recognition rate compared to existing methods and can be integrated into real time face recognition system.

Keywords

Face Recognition, Image Synthesis, Illumination Restoration, Dimensionality Reduction.

1. INTRODUCTION

Face Recognition has been recently receiving particular attention especially in security related fields. As for the early researches, both geometric feature based methods and template-matching methods were regarded as typical technologies. Since the 1990s appearance-based methods have been playing a dominant role in the area.

However, face recognition remains a difficult, unsolved problem in general. The changes induced by illumination are often larger than the differences between individuals, causing systems based directly on comparing images to misclassify input images. Approaches for handling variable illumination can be divided into four main categories: (1) extraction of illumination invariant features [Adi97, Lai01] (2) transformation of images with variable illuminations to a canonical representation [Xie05, Liu05, Sha01] (3) modeling the illumination variations [Geo01] (4) utilization of some 3D face models whose facial shapes and albedos are obtained in advance [Wey04, Zhan05].

Another effort related to varying illumination is the creation or synthesis of the image space of a novel image illuminated from an arbitrary angle. The works done on this topic consider face images as Lambertian surfaces, and at the same time assume faces as objects having the same shape but different surface texture [Sha01, Geo01, Geo03, Zhan05, Zhao00, Zhao03]. While Georghiades et al. [Geo01] do not deal with cast and attached shadows, Zhang et al. [Zhan05] minimize shadow effects by applying surface reconstruction.

In this work we propose a novel approach to minimize the effects of illumination variations on face recognition performance. Our method requires only one training image for any subject that will undergo testing. We try to solve the Small Sample Size (SSS) problem regarding this case. Our pipeline consists of three main steps: (1) synthesis of the image space for any input image, (2) training using a class-based linear discriminant approach, and (3) illumination restoration [Liu05] of any incoming face image during the recognition process. In Section 2 we will give a short introduction to PCA and LDA as dimensionality reduction techniques. An image synthesis method will be introduced in section 3. Section 4 will present an iterative process of illumination restoration. Finally experimental results will be given in section 5 and conclusions are drawn in Section 6.

2. DIMENSIONALITY REDUCTION

When using appearance based methods, we usually represent an image of width w and height h as a vector in a $w \cdot h$ dimensional space. In practice, this space, i.e. the full image space, is too large to allow robust and fast object recognition. A common way to attempt to resolve this problem is to use dimensionality reduction techniques.

2.1 Principal Component Analysis

Principal Component Analysis [Tur91] is a technique that is useful for the compression and classification of data. More formally let us consider a set of N sample images $\{p_1, p_2, \ldots, p_N\}$ taking values in an n-dimensional image space, and assume that each image belongs to one of c classes $\{P_1, P_2, \ldots, P_c\}$. We consider a linear transformation mapping the original n-dimensional image space into an m-dimensional feature space, where (m < n). If the total scatter matrix S_T (covariance matrix) is defined as:

$$S_T = \sum_{k=1}^{N} (p_k - \mu)(p_k - \mu)^T \qquad (1)$$

the dimension reduction is realized by solving the eigenvalues problem:

$$S_T \Phi = \Phi \Lambda \qquad (2)$$

where μ is the mean image, Λ is a diagonal matrix whose diagonal elements are the eigenvalues of S_T with their magnitudes in descending order, and Φ is a matrix whose i^{th} column is the i^{th} eigenvector of S_T. In order to obtain the eigenspace we generally choose m eigenvectors corresponding to the m largest eigenvalues, which capture over 95% of the variations in the images. After calculating the eigenfaces the projection is the only process left to be done. Let W_S be the matrix consisting of the m eigenvectors, and I_n be a new face image. The projection of I_n onto the eigenspace is represented as follows:

$$a = W_S^T (I_n - \mu) \qquad (3)$$

where a is mx1 vector containing the m projection coefficients. The reconstructed image is then given as:

$$I_n{}' = W_S a + \mu . \qquad (4)$$

The reconstructed image is the best approximation of the input image in the mean square sense. An advantage of using such representations is their reduced sensitivity to noise. A drawback of this approach is that the scatter being maximized is due not only to the between-class scatter that is useful for classification, but also to the within-class scatter that, for classification purposes, is unwanted information. Thus if PCA is presented with images of faces under varying illumination, the projection matrix W_S will contain principal components which retain, in the projected feature space, the variation due to lighting. Consequently, the points in the projected space will not be well clustered, and worse, the classes may be smeared together. It has been suggested that by discarding the three most significant principal components, the variation due to lighting is reduced [Bel97].

2.2 Fisher Linear Discriminant Analysis

Fisher's Linear Discriminant Analysis (FLDA) [Bel97] uses an important fact of photometric stereo: In the absence of shadowing, given three images of a Lambertian surface taken from the same viewpoint under three known, linearly independent light source directions, the albedo and surface normal can be recovered. For classification, this fact has great importance. It shows that, for a fixed viewpoint, the images of a Lambertian surface lie in a 3D linear subspace of the high-dimensional image space. One can perform dimensionality reduction using linear projection and still preserve linear separability. More formally, let us continue our previous study of PCA from this new perspective. The Linear Discriminant Analysis (LDA) selects W in such a way that the ratio of the between-class scatter and the within-class scatter is maximized. Let the between-class scatter matrix be defined as:

$$S_B = \sum_{i=1}^{C} N_i (\mu_i - \mu)(\mu_i - \mu)^T \qquad (5)$$

and the within-class scatter be defined as:

$$S_W = \sum_{i=1}^{C} \sum_{k=1}^{N} (p_k - \mu_i)(p_k - \mu_i)^T \qquad (6)$$

where μ_i is the mean image of class P_i, and N_i is the number of samples in class P_i. If S_W is nonsingular, the optimal projection W_{opt} is chosen as the matrix with orthonormal columns which maximizes the ratio of the determinant of the between-class scatter matrix of the projected samples to the determinant of the within-class scatter matrix of the projected samples, i.e.:

$$W_{opt} = \arg \max_{W} \frac{\left| W^T S_B W \right|}{\left| W^T S_W W \right|} = [w_1 \, w_2 \ldots w_m] \qquad (7)$$

In the face recognition problem, one is confronted with the difficulty that the within-class scatter matrix S_W is always singular. This stems from the fact that the rank of S_W is at most N - c, and, in general, the number of images in the learning set N is much smaller than the number of pixels in each image n. This means that it is possible to choose the matrix W such that the within-class scatter of the projected samples can be made exactly zero. To overcome the complication of a singular S_W, an alternative method has been proposed called Fisherfaces. It avoids this

problem by projecting the image set to a lower dimensional space so that the resulting within-class scatter matrix S_W is nonsingular. This is achieved by using PCA to reduce the dimension of the feature space to N - c, and then applying the standard LDA defined by Eq. (7) to reduce the dimension to c - 1.

Recently it has been shown that the null space of S_W may contain significant discriminatory information [Lu03, Gao06]. As a consequence, some of the significant discriminatory information may be lost due to the preprocessing PCA step. Many methods classified as Direct LDA [Chen00, Yu00] have been developed to deal with this problem. However, the Fisherface method appears to be the best simultaneously handling variation in lighting. It has lower error rate than the PCA method.

3. IMAGE SYNTHESIS

Nearly all approaches to view synthesis take a set of images gathered from multiple viewpoints and apply techniques related to structure from motion, stereopsis, image transfer, image warping, or image morphing. Each of these methods requires the establishment of correspondences between image data (e.g., pixels) across the set. Since dense correspondence is difficult to obtain, most methods extract sparse image features (e.g., corners, lines), and may use multi-view geometric constraints or scene-dependent geometric constraints to reduce the search process and constrain the estimates. For these approaches to be effective, there must be sufficient texture or viewpoint-independent scene features, such as albedo discontinuities or surface normal discontinuities. Underlying nearly all such stereo algorithms is a constant brightness assumption, that is, the intensity (irradiance) of corresponding pixels should be the same.

This section is based on a work showing that the set of all images generated by varying lighting conditions on a collection of Lambertian objects can be characterized analytically using images of a prototype object and a (illumination invariant) "signature" image per object of the class called Quotient Image [Sha01]. In this approach the consideration will be restricted to objects with a Lambertian reflectance:

$$I(k,l) = \rho(k,l)n(k,l)^T s \qquad (8)$$

where $0 \leq \rho(k, l) \leq 1$ is the surface texture, n(k, l) is the surface normal direction, and s is the light source direction whose magnitude is the light source intensity. Furthermore it is assumed that the faces belonging to a class have the same shape but differ in surface texture. Although this is a very strong assumption it can be shown that it holds if faces are roughly aligned.

Given two objects y and a, let the quotient image Q be the ratio of their albedos:

$$Q_y(u,v) = \frac{\rho_y(u,v)}{\rho_a(u,v)} \qquad (9)$$

where u, v change over the image. Clearly, Q is illumination invariant. The importance of this ratio becomes clear by the following statement:

Given three images a_1, a_2, a_3 of object a, illuminated by any three linearly independent lighting conditions and an image y_s of y illuminated by some light sources, then there exists coefficient x_1, x_2, x_3 that satisfy:

$$y_S = (\sum_j x_j a_j) \otimes Q_y \qquad (10)$$

where $\otimes$ denotes the Cartesian product (pixel by pixel multiplication).

We see that once Q_y is given, we can generate y_s (the novel image) and all other images of the image space of y. The key is to obtain the correct coefficients x_j which can be done by using a bootstrap. Let the bootstrap set of 3N pictures be taken from three fixed (linearly independent) not necessarily known light sources s_1, s_2 and s_3. Let A_i, $i = 1,..., N$ be a matrix whose columns are the three pictures of object a_i with albedo function ρ_i. Thus, $A_1,...,A_N$ represent the bootstrap set of N matrices, each is a (nx3) matrix, where n is the number of pixels of the image. Let y_s be an image of some novel object y (not part of the bootstrap set) illuminated by some light source $s=\sum_j x_j s_j$. We wish to recover $x=\{x_1,x_2,x_3\}$ given the N matrices $A_1,..., A_N$ and the vector y_s. This can be done by solving a bilinear problem in the N+3 unknowns x and α_i, which can be obtained by minimizing the function:

$$f(x) = \frac{1}{2} \sum_{i=1}^{N} \left| A_i x - \alpha_i y_s \right|^2 \qquad (11)$$

for the unknown x. To find the desired global minimum we apply the Euler-Lagrange equation related with the variables x and α. This can be done by derivation of $f(x)$ through these variables. We get the following relations:

$$x = \sum_{i=1}^{N} \alpha_i v_i \qquad (12)$$

Fig. 1. Example of the image synthesis. The input image is in the upper leftmost position

$$v_i = \left(\sum_{r=1}^{N} A_r^T A_r \right)^{-1} A_i^T y_s \qquad (13)$$

$$\alpha_i y_S^T y_S - \left(\sum_{r=1}^{N} \alpha_r v_r \right)^T A_i^T y_S = 0 \qquad (14)$$

So we first find the v vectors (3x1) in Eq. (13), and we solve the homogeneous linear equation in Eq. (14) for the α_i. Then by using Eq. (12) the desired minimum can be found. Finally we compute the quotient image $Q_y = y_s / Ax$, where A is the average of $A_1, ..., A_N$. The image space is spanned by the product of images Q_y and Az for all choices of z. An example of an output based on a training bootstrap (10 persons) from YaleB is given in Fig 1.

As we see, the results of this algorithm are quite satisfactory in spanning the image space of a given input image. This helps to overcome the SSS problem and creates the possibility to use LDA-based methods even when only one image per person is provided during the learning phase. An inherent assumption throughout the algorithm is that for a given pixel (x, y), n(x, y) is the same for all the images, i.e, the bootstrap set as well as the test images. The performance degrades when dominant features between the bootstrap set and the test are misaligned. As this step will occur after the face detection it is supposed that dominant features will have been depicted and aligned previously.

4. ITERATIVE METHOD IN ILLUMINATION RESTORATION

The major advantages of the algorithm explained in this section are that no facial feature extraction is needed and the generated face images will be visually natural looking. The method is based on the general idea that the ratio of two images of the same

person is simpler to deal with than directly comparing images of different persons.

It uses a ratio-image, which is the quotient between a face image whose lighting condition is to be normalized and a reference face image. The two images are blurred using a Gaussian filter, and the reference image is then updated by an iterative strategy in order to further improve the quality of the restored face image. In this approach, a face image with arbitrary illumination can be restored so as to have frontal illumination.

Let I_{ik} denote a face image of the ith person captured under the s_k light source direction, where a light source is classified according to its direction. I_{r0} represents a face image of another person captured under the frontal light source s_0 and is used as a reference image. Then, we give the two blurred images of I_{ik} and I_{r0} denoted as B_{ik} and B_{r0}, respectively, as:

$$B_{ik} = F * I_{ik} = F * (\rho_i n_i^T s_k) = (F * \rho_i n_i^T) s_k \quad (15)$$

$$B_{r0} = F * I_{r0} = F * (\rho_r n_r^T s_0) = (F * \rho_r n_r^T) s_0 \quad (16)$$

where (*) is the convolution operation and F is a 2D Gaussian low-pass filter, with $\sigma_x = \sigma_y = \sigma$, given by the following formula:

$$F(x, y) = \frac{1}{2\pi\sigma^2} e^{-(x^2 + y^2)/2\sigma^2} \qquad (17)$$

As the shape and albedos of all faces are similar if the size of F is big enough, we can assume that $B_{i0} \approx B_{r0}$. By using the formulas (15)-(17) and this assumption, we can obtain the face image under frontal illumination for the ith person from I_{ik} captured under an arbitrary lighting direction s_k by:

$$H_{io} = \rho_i n_i^T s_o \approx \rho_i n_i^T s_k \frac{\left(F * \rho_r n_r^T\right) s_o}{\left(F * \rho_i n_i^T\right) s_k} = I_{ik} \frac{B_{ro}}{B_{ik}} \quad (18)$$

H_{i0} is an estimation of I_{i0} or a restored image of I_{ik}. This approach can be summarized by the following algorithm:

1. A mean face image and an eigenspace Φ_S introduced in section 2 are computed based on a set of training images, which are all captured under frontal light source

2. An initial restored image can be calculated using Eq. (15)-(18), where the mean face image is used as the initial reference image and the size of the Gaussian filter F is 5. A reconstructed image is obtained from the initial restored image based on the computed eigenspace, and should have a smaller number of noise points.

3. An iterative procedure is used to obtain a final restored image with frontal illumination. During the iterative procedure, the reference image is updated with the new reconstructed image so as to obtain a visually better restored image.

4. The iterative procedure continues until a stopping criterion is met. In this approach the stopping criterion is the difference between two consecutive outputs of Eq. (18), or a specified maximum number of iterations.

5. EXPERIMENTAL RESULTS

All the experiments have been done using the YaleB database. Despite its relatively small size, this database contains samples from the whole illumination space and has become a testing standard for variable-illumination recognition-methods. This database consists of 10 distinct persons and 45 images for each person, divided in four subsets. The first subset includes 70 images captured under light source directions within a span of 12˚. The second and third subsets have respectively 120 images and 140 images each, captured under light source directions within 25˚ and within 50˚. The fourth subset contains 120 images captured under light source directions within 75˚. In all the images the position of the two eyes of each face were located and translated to the same position. The images were cropped to a size of 180x190. In order to improve the performance of dimensionality reduction and recognition, all the images were normalized to zero mean and unit variance. After that the pipeline was tested when histogram equalization (HE) and adaptive histogram equalization (AHE) was applied

as a further preprocessing step. Table 1 shows the effect of the preprocessing on the recognition rate, where it is obvious that AHE gives the best result among these preprocessing techniques.

	No	HE	AHE
Recognition	43.4	74	81.5

Table 1. Results with the YaleB database (PCA used)

A wide range of experiments have been conducted to test the Quotient Image algorithm. First, synthesizing new images from any arbitrarily illuminated input image outside the YaleB database is considered by using a bootstrap consisting of 30 pictures of 10 persons of YaleB (Fig 1). Furthermore we examined the performance when only 15 images (5 persons), 9 images (3 persons) and 3 images (1 person) respectively were available in the bootstrap (Fig 2). In all these cases, after calculating the Quotient Image by means of $Q_y=y_s/Ax,$ we create the image space of the novel image by the product of Q_y and Az for all choices of z. Some examples of the calculated **x** variable are given in Table 2.

Coeff/#Person	5	3	1
x_1	0.11302	0.23729	0.15915
x_2	0.38648	0.31587	0.45989
x_3	0.41526	0.35312	0.29723

Table 2. Coefficient results for different bootstrap combinations

Using these coefficients, we create the image space of the input image by randomly assigning different values for z or using a normal distribution around the original values of *X*. From Fig. 2 we can see that a bootstrap consisting of 10 persons is quite consistent for creating the image space of an input image. Even when we reduce it by half, the results are quite satisfactory. This is because the albedos of possible faces occupy only a small part of the dimension in which they are spanned. Of course the larger the bootstrap size the more accurate will be the recovery of **x** and the quotient image.

In order to prepare the training set for the LDA process we create the image space of all 10 persons of the YaleB database. For these we used 15 images for bootstrap where the object being reconstructed has been left out. The results of the LDA step are given in Table 3 and Table 4. The final step of our approach is to reconstruct any incoming image in order to have a frontal illuminated image. In the experiments with the YaleB database subsets the results were almost identical to the frontal illuminated image for the first and second subset.

(a)

(b)

(c)

(d)

Fig. 2. Image re-rendering results for different bootstrap combinations: (a) 10 persons (b) 5 persons; (c) 3 persons; (d) 1 person

The other two subsets where the illumination conditions are worse also produce high performance but it has to be stated that some noise and feature corruption became visible (Fig 3).

After the illumination restoration process the performance of the whole approach was tested with 450 input images. Different distance measurements were experimented with:

- Manhattan distance (L1- norm)
- Cosine angle between two vector representations
- Euclidian distance (L2 - norm)

The Euclidian distance gives the best results for this classification purpose when used in a one-nearest neighbor classifier (Table 3, 4).

In order to further increase the recognition rate several combinations during the training phase have been applied. For the LDA step the best performance was achieved when 10 synthesized images per object

were available during the training phase and all the discriminatory feature vectors of the LDA-projection matrix were used (feature vectors of length 9). The use of a higher number of synthesized images slightly increases the performance. In all the experiments the results were compared with PCA because it is the most important discriminatory technique used when only one image per person is available during training.

	Subset1	Subset2	Subset3	Subset4	Total
HE+PCA (%)	100	97.5	66.42	44.16	74
HE+New (%)	100	100	92.8	91.66	95.56

Table 3. Recognition rates with histogram equalization preprocessing

	Subset1	Subset2	Subset3	Subset4	Total
AHE+PCA	100	100	83.57	50	81.55
AHE+New	100	100	100	95	98.6

Table 4. Recognition rates with adaptive histogram equalization preprocessing

(a)

(b)

Fig. 3. Illumination restoration for images of Subset4 (up to 70∘) (a) Before preprocessing and illumination restoration (b) After illumination restoration

Obviously our proposed approach can significantly improve the recognition rates. Other methods have achieved high recognition rates on the YaleB database, but they require a large number of images for each person. A recent work [Zhan05] proposes an illumination invariant face recognition system from a single input image. It tries to solve the problem by using spherical-harmonics basis-images and they achieve very good results. However they specify that the computational burden of such an algorithm is high for the integration in a real time system.

6. CONCLUSION

A novel pipeline for dealing with illumination variation was introduced. In this work was aimed at a solution of the SSS problem of class-based discrimination problems. For this an image-space synthesis-method was explained and the image space of each image of the training set was created. After creating the image space of each training image, FLDA was applied in order to best use the discriminatory features of the system.

Every incoming image was processed with the illumination restoration algorithm and then a projection was done in order to extract the discriminatory features. The recognition rate with the YaleB database consisting of 450 images was 98.66% which can be considered very successful when compared to existing methods. Another approach [Geo01a] claimed 100% recognition rates in all data subsets, but seven images of each person in Subset1 have to be used to obtain the shape and albedo of the face.

As a conclusion, this work proposes an innovative approach for creating a robust face recognition system under varying illumination. This study offers the possibility of creating a real time system because it is not computationally complex. In the future this study will be extended to deal not only with upright frontal views but also with different poses. One possible approach based on this study is to apply multiple reference subspaces for different poses.

7. REFERENCES

[Adi97] Adini, Y., Moses, Y., and Ullman, S.: Face recognition: The problem of compensating for changes in illumination direction. PAMI, Vol. 19, No. 7, pp. 721–732, 1997.

[Bel97] Belhumeur, P.N, Hespanha, J.P., and Kriegman, D.J.: Eigenfaces vs. Fisherfaces: Recognition using class specific linear projection. IEEE Trans. PAMI, Vol. 19, No. 7, pp 711-720, 1997.

[Chen00] Chen, L., Liao, H.M., Ko, M., Lin, J., Yu, G.: A new LDA-based face recognition system which can solve the small sample size problem. Pattern Recognition Letters, Vol. 33, pp 1713-1726, 2000.

[Gao06] Gao, H., Davis, J.W.: Why direct LDA is not equivalent to LDA. Pattern Recognition Letters, Vol. 39, pp 1002-1006, 2006.

[Geo01] Georghiades, A.S., Belhumeur, P.N., Kriegman, D.J.: From few to many: Illumination cone models for face recognition under variable lighting and pose. IEEE Trans. Pattern Anal. Mach. Intell. 23 (2) pp. 643–660, 2001.

[Geo03] Georghiades, A.S., Belhumeur, P.N., Kriegman, D.J.: Illumination-based image synthesis: Creating novel images of human faces under differing pose and lighting. IEEE workshop on Multi-View Modeling and Analysis of Visual Scenes, pp 47 - 54, 1999.

[Lai01] Lai, J.H., Yuen, P.C., Feng, G.C.: Face recognition using holistic Fourier invariant features. Pattern Recognition Letters, Vol. 34, pp 95-109, 2001.

[Liu05] Liu, D.H., Lam, K. M., Shen, L. S.: Illumination invariant face recognition. Pattern Recognition, Vol. 38, No 10, pp. 1705-1716, 2005.

[Lu03] Lu, J., Plataniotis, K.N., Venetsanopoulos, A.N.: Regularization studies of linear discriminant analysis in small sample size scenarios with application to face recognition. Pattern Recognition Letters, Vol. 26, No. 2, pp 181-191, 2005.

[Sha01] Shashua, A., Tammy, R.R.: The quotient image: Class-based re-rendering and recognition with varying illuminations. IEEE Trans. Pattern Anal. Mach. Intell. 23 (2), pp. 129–139, 2001.

[Tur91] Turk, M., Pentland, A.: Eigenfaces for recognition. Journal of Cognitive Neuroscience, Vol. 3, Num. 1, pp. 71-86, 1991.

[Wey04] Weyrauch, B., Heisele, B., Huang, J., Blanz, V.: Component-Based face recognition with 3D morphable models. IEEE Trans. CVPRW'04, Vol. 5, pp 85-90, 2004.

[Xie05] Xie, X., Lam, K.M.: Face recognition under varying illumination based on a 2D face shape model. Pattern Recognition, Vol.38, No 2, pp. 221-230, 2005.

[Yu00] Yu, H., Yang, J.: A direct LDA algorithm for high-dimensional data-with application to face recognition. Pattern Recognition Letters, Vol. 34, pp 2067-2070, 2000.

[Zhan05] Zhang, L., Wang, S., Samaras, D.: Face synthesis and recognition from a single image under arbitrary unknown lighting using a spherical harmonic basis morphable model. CVPR 2005, Volume 2, pp. 209 – 216, 2005.

[Zhao00] Zhao, W., Chellapa, R: SFS based view synthesis for robust face recognition. Proceedings of the 4th Conference on Automatic Face and Gesture Recognition, pp 285-292, 2000.

[Zhao03] Zhao, W., Chellapa, R., Philips, P.J., Rosenfeld, A.: Face Recognition: A literature Survey. ACM Computing Surveys, Vol. 35, No. 4, pp. 399-458, 2003.