

Pose Invariant Face Recognition under Arbitrary Unknown Lighting using Spherical Harmonics

Lei Zhang and Dimitris Samaras

Department of Computer Science,
SUNY at Stony Brook, NY, 11790
{lzhang, samaras}@cs.sunysb.edu

Abstract. We propose a new method for face recognition under arbitrary pose and illumination conditions, which requires only one training image per subject. Furthermore, no limitation on the pose and illumination conditions for the training image is necessary. Our method combines the strengths of Morphable models to capture the variability of 3D face shape and a spherical harmonic representation for the illumination. Morphable models are successful in 3D face reconstructions from one single image. Recent research demonstrates that the set of images of a convex Lambertian object obtained under a wide variety of lighting conditions can be approximated accurately by a low-dimensional linear subspace using spherical harmonics representation. In this paper, we show that we can recover the 3D faces with texture information from one single training image under arbitrary illumination conditions and perform robust pose and illumination invariant face recognition by using the recovered 3D faces. During training, given an image under arbitrary illumination, we first compute the shape parameters from a shape error estimated by the displacements of a set of feature points. Then we estimate the illumination coefficients and texture information using the spherical harmonics illumination representation. The reconstructed 3D models serve as generative models to render sets of basis images of each subject for different poses. During testing, we recognize the face for which there exists a weighted combination of basis images that is the closest to the test face image. We provide a series of experiments on approximately 5000 images from the CMU-PIE database. We achieve high recognition rates for images under a wide range of illumination conditions, including multiple sources of illumination.

1 Introduction

Face recognition has recently received extensive attention as one of the most significant applications of image understanding. Although rapid progress has been made in this area during the last few years [29][21][16][3][35][5][19][18][8][28][24][9][20][31][33], the general task of recognition remains unsolved. In general, face appearance does not depend solely on identity. It is also influenced by illumination and viewpoint. Changes in pose and illumination will cause large changes in the appearance of a face. In this paper we demonstrate a new method to recognize face images under a wide range of pose and illumination conditions using spherical harmonic images of the face and a morphable model. Our method requires only a single training image per subject. To our knowledge no other face recognition method can achieve such a high level of pose and illumination invariance when only one training image is available.

In the past few years, there have been attempts to address image variation produced by changing in illumination and pose [10][35]. Georghiades et al. [11] present a new method using the illumination cone which requires at least three images per subject to build the illumination cone. Romdhani et al. [23] recover the shape and texture parameters of a 3D Morphable Model in an analysis-by-synthesis fashion. In [23], the shape parameters are computed from a shape error estimated by optical flow and the texture parameters are obtained from a texture error. The algorithm uses linear equations to recover the shape and texture parameters irrespective of pose and lighting conditions of the face image. However, this method is bound to images taken under single directional illumination and requires the knowledge of light direction which is difficult to know in most cases.

In general, appearance-based methods like Eigenfaces [29] and SLAM [21] need a number of training images for each subject, in order to cope with pose and illumination variability. Previous research suggests that illumination variability in face images is low-dimensional e.g. [2][22][12][4][1][25][26][17]. Using spherical harmonics and signal-processing techniques, Basri et al. [2] and Ramamoorthi [22] have independently shown that the set of images of a convex Lambertian object obtained under a wide variety of lighting conditions can be approximated accurately by a 9 dimensional linear subspace. Furthermore, a simple scheme for face recognition with excellent results was described in [2]. However, to use this recognition scheme, the basis images spanning the illumination space for each face are required. These images can be rendered from a 3D scan of the face or can be estimated by applying PCA to a number of images of the same subject under different illuminations [22]. An effective approximation of this basis by 9 single light source images of a face was reported in [15] and Wang et al. [30] proposed an illumination modeling and normalization method for face recognition. The above mentioned methods need a number of training images and/or 3D scans of the subjects in the database, requiring specialized equipment and procedures for the capture of the training set, thus limiting their applicability. A promising earlier attempt by [36] used symmetric shape from shading but suffered from the drawbacks of SFS. A new approach is proposed in [34] for face recognition under

arbitrary illumination conditions, for fixed pose, which requires only one training image per subject and no 3D shape information. In [34] the statistical model is based on a collection of 2D basis images, rendered from known 3D shapes. Thus 3D shape is only implicitly included in the statistical model. Here we will base our statistical model directly on 3D shapes, perform statistical analysis in 3D in order to estimate the most appropriate 3D shape and then create the 2D basis images. The ability to manipulate the 3D shape explicitly allows the generation of basis images for poses that do not exist in the training data.

In this paper we propose a method that combines a 3D morphable model and a low-dimensional illumination representation that uses spherical harmonics. Our method requires only one training image for each subject without pose and illumination limitations. Our method consists of three steps: 3D face reconstruction, basis image rendering and recognition. Initially, similar to [23], given a training image, we compute the shape parameters of a morphable model from a shape error estimated by the displacements of a set of feature points. Then we estimate the illumination coefficients and texture information using the spherical harmonics illumination representation. In the basis image rendering step, the reconstructed face models then serve as generative models that can be used to synthesize sets of basis images under novel poses and spanning the illumination field. During the recognition step, we use the recognition scheme proposed by Basri et al. [2]. We return the face from the training set for which there exists a weighted combination of basis images that is the closest to the test face image.

We use the morphable model computed from the USF 3D face database [6] and the CMU-PIE database [27] for training and testing. We provide a series of experiments that show that the method achieves high recognition accuracy although our method requires only a single training image without limitation on pose and illumination conditions. We compare the recognition rate with [23] on the images taken under single light source. We also give experiment results of recognition on the set of images under multiple light sources, and compare with [34] for known pose.

This paper is organized as follows. In the next section, we will briefly introduce the Morphable Model. In Section 3, we explain the Spherical Harmonics and how to acquire basis images from 3D face models. In Section 4, we describe the process of 3D face model reconstruction and basis image rendering. In Section 5, we describe the recognition process that uses the rendered basis images. In Section 6, we describe our experiments and their results. The final Section presents the conclusions and future work directions.

2 Morphable Model

In this section we briefly summarize the morphable model framework described in detail in [6][7]. The 3D Morphable Face Model is a 3D model of faces with separate shape and texture models that are learnt from a set of exemplar faces. Morphing between faces requires complete sets of correspondences between all of the faces. When building a 3D morphable model, we transform the shape and texture spaces into vector spaces, so that any convex combination of exemplar

shapes and textures represents a realistic human face. We present the geometry of a face with a shape-vector $S = (X_1, Y_1, Z_1, X_2, \dots, Y_n, Z_n)^T \in \mathbb{R}^{3n}$, which contains the X, Y, Z -coordinates of its n vertices. Similarly, the texture of a face can be represented by a texture-vector $T = (R_1, G_1, B_1, R_2, \dots, G_n, B_n)^T \in \mathbb{R}^{3n}$ where the R, G, B texture values are sampled at the same n points. A morphable model can be constructed using a data set of m exemplar faces; exemplar i is represented by the shape-vector S_i and texture-vector T_i . New shapes s and textures t can be generated by convex combinations of the shapes and textures of the m exemplar faces: $s = \sum_{i=1}^m a_i S_i, t = \sum_{i=1}^m b_i T_i, \sum_{i=1}^m a_i = \sum_{i=1}^m b_i = 1$. To reduce the dimensionality of the shape and texture spaces, Principal Component Analysis (PCA) is applied separately on the shape and texture spaces:

$$s = \bar{s} + \sum_{i=1}^{m-1} \alpha_i \sigma_{s,i} s_i, \quad t = \bar{t} + \sum_{i=1}^{m-1} \beta_i \sigma_{t,i} t_i \quad (1)$$

By setting the smallest eigenvalues to zero, Eq. 1 is reformulated as:

$$s = \bar{s} + S\alpha, \quad t = \bar{t} + T\beta \quad (2)$$

In Eq. 2 the columns of S and T are the most significant eigenvectors s_i and t_i re-scaled by their standard deviation and the coefficients α and β constitute a pose and illumination invariant low-dimensional coding of a face [23]. PCA also provides an estimate of the probability densities of the shapes and textures, under a Gaussian assumption: $p(s) \sim e^{-\frac{1}{2}\|\alpha\|^2}$, $p(t) \sim e^{-\frac{1}{2}\|\beta\|^2}$

3 Spherical Harmonics

In this section, we will briefly explain the illumination representation by using spherical harmonics and how we render basis images from 3D models using the results of [2]. Let L denote the distant lighting distribution. By neglecting the cast shadows and near-field illumination, the irradiance E is then a function of the surface normal n only and is given by an integral over the upper hemisphere Ω_n [22]: $E(n) = \int L(\omega)(n \cdot \omega)d\omega$ We then scale E by the surface albedo λ to find the radiosity I , which corresponds to the image intensity directly:

$$I(p, n) = \lambda(p)E(n) \quad (3)$$

Basri et al. [2] and Ramamoorthi [22] have independently shown that E can be approximated by the combination of the first nine spherical harmonics $H(x, y, z)$ for Lambertian surfaces:

$$\begin{aligned} h_{00} &= \frac{1}{\sqrt{4\pi}}, & h_{11}^o &= \sqrt{\frac{3}{4\pi}}y, & h_{21}^o &= 3\sqrt{\frac{5}{12\pi}}yz \\ h_{10} &= \sqrt{\frac{3}{4\pi}}z, & h_{20} &= \frac{1}{2}\sqrt{\frac{5}{4\pi}}(2z^2 - x^2 - y^2), & h_{22}^e &= \frac{3}{2}\sqrt{\frac{5}{12\pi}}(x^2 - y^2) \\ h_{11}^e &= \sqrt{\frac{3}{4\pi}}x, & h_{21}^e &= 3\sqrt{\frac{5}{12\pi}}xz, & h_{22}^o &= 3\sqrt{\frac{5}{12\pi}}xy \end{aligned} \quad (4)$$

where the superscripts e and o denote the even and the odd components of the harmonics respectively and x, y, z demote the cartesian components. Then the image intensity of a point p with surface normal $n = (n_x, n_y, n_z)$ and albedo λ can be computed according to Eq. 3 by replacing x, y, z with n_x, n_y, n_z . Fig. 1 gives an example of the mean shape and texture of the morphable model under a spherical harmonics representation.

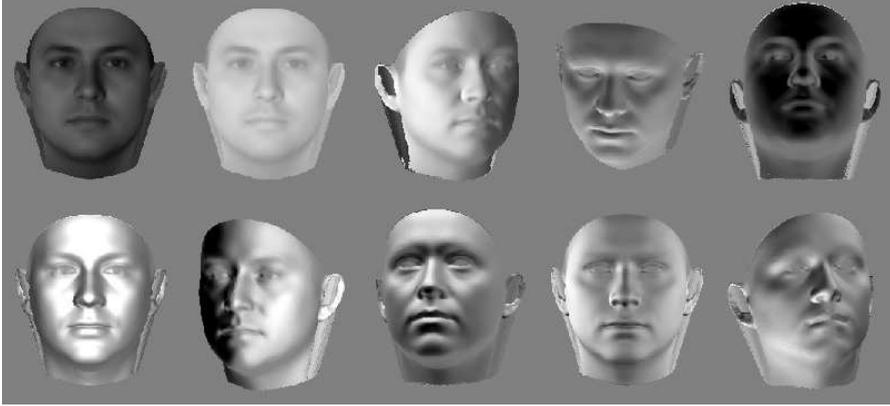


Fig. 1. The first image is the mean of the morphable model and the following nine images are the basis images under various view-points, represented by spherical harmonics. Lighter gray (0-127) represents positive values and darker gray (128-255) represents negative values.

4 Face Model Reconstruction and Basis Image Rendering

In this section, we will explain how we recover the shape and texture information of a training subject by combining a morphable model and spherical harmonics lighting representation.

4.1 Forward and Inverse Face Rendering

We can generate photo-realistic face images by using the morphable model we described in Section 2 [6]. Here we describe how we synthesize a new face image from the face shape and texture vectors s and t , thus, the inversion process of the synthesis is how we recover shape and texture information from the image.

Shape: Similar to [23], a realistic face shape can be generated by:

$$s_{2d} = fPR(\bar{s} + S\alpha + t_{3d}) + t_{2d} \quad (5)$$

where f is a scale parameter, P an orthographic projection matrix and R a rotation matrix with ϕ , γ and θ the three rotation angles for the three axes. t_{3d} and t_{2d} are translation vectors in 3D and 2D respectively. Eq. 5 relates the vector of 2D image coordinates s_{2d} and the shape parameters α . For rendering, a visibility test must still be performed by using a z -buffer method.

For a training image, inverting the rendering process, the shape parameters can be recovered from the shape error: if f , ϕ , γ and θ are kept constant, the relation between the shape s_{2d} and α is linear according to Eq. 5: $\frac{\partial s_{2d}}{\partial \alpha} = fPRS$. Thus, updating α from a shape error δs_{2d} requires only the solution of a linear system of equations. In our method, the shape error is estimated by the displacements of a set of manually picked feature points s_f [14] corresponding to image coordinates s_f^{img} . The shape reconstruction goes through the following steps:

Model Initialization: All the parameters are initialized in this step. Shape parameter α is set to 0 and pose parameters f, ϕ, γ, θ and t_{2d} are initialized manually. We do not need to know the illumination conditions of the training image, unlike [23].

Feature Correspondence: For the set of pre-picked feature points in the morphable model, we find the correspondence s_f^{img} in the training image semi-automatically. The set of feature points contains major and secondary features, see Fig. 2. After the correspondences of major features are manually set, the secondary features are updated automatically.

Rotation, Translation and Scale Parameters Update: the parameters f , ϕ , γ and θ can be recovered by using a Levenberg-Marquardt optimization to minimize the error between s_f^{img} and the model feature points [13]:

$$\operatorname{argmin}_{f, \phi, \gamma, \theta, t_{2d}} \|s_f^{img} - (fPR(\bar{s}_f + S_f\alpha + t_{3d}) + t_{2d})\|^2 = (\tilde{f}, \tilde{\phi}, \tilde{\gamma}, \tilde{\theta}, \tilde{t}_{2d}) \quad (6)$$

where \bar{s}_f and S_f is the corresponding shape information of the feature points in the morphable model in Eq. 2.

Shape Parameter Update: The shape error of the feature points, δs_f^{2d} , is defined as the difference between s_f^{img} and the new shape information of feature points in the model that was rendered by recovered parameters $\tilde{f}, \tilde{\phi}, \tilde{\gamma}, \tilde{\theta}$ and \tilde{t}_{2d} . Thus, the vector of shape parameters α can be updated by solving a linear system of equations:

$$\delta s_f^{2d} = fPRS_f\delta\alpha \quad (7)$$

Texture: For texture information recovery, most of the previous methods [11][23] are applicable to images taken under single light source, which limits their applicability. Here we propose a method which performs texture fitting to a training image and has no limitation in the image illumination conditions.

According to Eq. 3 and 4, the texture of a face can be generated by:

$$t = B * l, \quad B = H(n_x, n_y, n_z) \cdot \lambda \quad (8)$$

where H is the spherical harmonics representation of the reflectance function (Eq. 4) and l is the vector of illumination coefficients. Hence, if we know the illumination coefficients, the texture information is only dependent on image intensity t and surface normal n , which can be computed from the 3D shape we recovered during the shape fitting step. The texture recovery is described as following:

Basis Computation: The initial albedo λ for each vertex is set to \bar{t} . With the recovered shape information, we first compute the surface normal n for each



Fig. 2. Recovery Results: Images in the first row are the input training images, those in the second row are the initial fittings, the third row shows images of the recovered 3D face model and the last row gives the illuminated rotated face models. In the first column, the black points are pre-picked major features, the white points are the corresponding features and the points lying in the white line are secondary features.

vertex. Then the first nine basis images B and spherical harmonics $H(n)$ for reflectance function can be computed according to Eq. 8 and 4 respectively.

Illumination Coefficients Estimation: The set of illumination coefficients l is updated by solving a linear system of equations:

$$t_{tra} = B_{cur}l \quad (9)$$

Texture Recovery: According to Eq. 8, the texture λ for each visible vertex is computed by solving: $t_{tra} = H(n_x, n_y, n_z)l \cdot \lambda$. Since texture is dependent on both current texture and illumination coefficients, the new value of λ is:

$$\lambda = (1 - \eta)\lambda_{cur} + \eta(t_{tra}/(H(n_x, n_y, n_z)l)) \quad (10)$$

We compute l and λ by solving Eq. 9 and 10 iteratively. In our experiments, weight η is first set to 0.5, then incremented by 0.1 at each step until it reaches 1. Instead of recovering texture parameters [23], we estimate the albedo value for each vertex, which will be used for basis image rendering and recognition. For occluded vertices, texture information is estimated through facial symmetry. Fig. 2 shows the results of our method.

There is human interactivity in our shape fitting part since we manually find the correspondences for major features. Automatic shape fitting of a morphable model [23] is beyond the scope of this paper which focuses on the statistics of interaction of geometry with arbitrary unknown illumination and the feature based method performed sufficiently well to demonstrate the strength of our approach.

4.2 Basis Images Rendering

For each training subject, we recover a 3D face model using the algorithm described in section 4.2. The recovered face models serve as generative models to render basis images. In this section, for each subject, a set of basis images across poses are generated, to be used during recognition. We sample the pose variance for each 5° in both azimuth and altitude axes. In our experiments, the range of azimuth is $[-70, 70]$ and the range of altitude is $[-10, 10]$. Fig. 3 shows a subset of the basis images for one subject.

5 Face Recognition

In the basis image rendering step, for each subject i , a set of 145 (29×5) basis $B_j^i, j \in [1..145]$ is rendered. During testing, given a new testing image I_t , we recognize the face of subject i for which there exists a weighted combination of basis images that is the closest to the test face image [2]: $\min_{i,j} \|B_j^i l - I_t\|$ where B_j^i is a set of basis images with size $d \times r$, d is the number of points in the image and r the number of basis images used (9 is a natural choice, we also tried 4 in our experiments). Every column of B_j^i contains one spherical harmonic image, and the columns of B_j^i form a basis for the linear subspace. To solve the equation, we simply apply QR decomposition to B_j^i to obtain an orthonormal basis. Thus, we compute the distance from the test image, I_t , and the space spanned by B_j^i as $\|QQ^T I_t - I_t\|$.



Fig. 3. A subset of the rendered basis images across poses.

Table 1. Recognition results and comparison: The first column lists the light numbers and the following two columns list the recognition rate for each pose. The recognition rates of the LiST algorithm are taken from [23].

Light	Front Gallery Using all 9 basis			Front Gallery Using first 4 basis			Side Gallery Using all 9 basis			Side Gallery Using first 4 basis		
	Front	Side	Profile	Front	Side	Profile	Front	Side	Profile	Front	Side	Profile
1	95	89	51	89	81	49	91	92	52	79	78	52
2	89	81	34	79	73	31	80	83	34	67	67	33
3	97	88	44	89	79	42	92	96	50	83	86	48
4	98	91	52	89	83	50	94	96	62	85	88	55
5	99	89	57	89	84	52	94	97	64	87	90	59
6	100	92	55	91	86	50	100	100	64	89	89	59
7	99	95	54	88	83	51	92	96	60	86	89	58
8	99	93	62	94	89	54	94	99	70	84	85	62
9	100	96	61	90	88	55	95	100	71	83	92	64
10	100	97	60	92	87	56	98	100	69	89	88	64
11	100	98	58	88	89	50	97	95	63	90	90	63
12	98	99	61	90	88	57	94	98	72	84	86	69
13	99	93	55	89	88	50	98	99	70	87	89	63
14	100	94	53	91	86	49	95	100	62	91	92	58
15	100	93	54	91	87	49	98	99	61	89	89	52
16	99	91	53	91	82	49	95	97	60	87	89	59
17	98	92	55	91	80	50	96	99	65	89	91	61
18	95	88	52	90	78	47	92	94	62	82	87	58
19	98	90	56	89	81	51	92	95	61	82	91	55
20	99	94	58	88	80	51	93	97	63	86	88	57
21	99	96	51	90	81	50	96	96	55	84	89	54
22	99	95	62	89	81	56	94	99	65	79	90	60
mean	98.2	92.4	54.5	89.4	83.4	50.0	94.4	96.7	61.6	84.6	87.4	57.4
LiST mean	97	91	60				93	96	71			

6 Experiments and Results

In our experiments, we used the CMU-PIE database which provides images of both pose and illumination variation. The CMU-PIE database contains 68 individuals, none of which is also in the USF set used to compute the morphable model. We performed experiments on a set of 4488 images which contains 68 subjects, 3 poses for each subject and 22 different illuminations for each pose.

6.1 Experiments of Illumination Invariance

Since our recognition method is based on the 3D face models recovered during training, it is important that the recovered face models and rendered basis images are robust. Figure 4 shows three sets of rendered basis images recovered from various face images under different illuminations for one subject. The resulting basis images rendered from images under different illumination are very close. For each subject, we calculated 10 sets of basis using 10 training images under different illumination. The per person mean variance of the 10 resulting sets of basis images was 3.32. For comparison, per person variance of the original training images was 20.25. That means the rendered basis images have much greater invariance to illumination effects than original images.

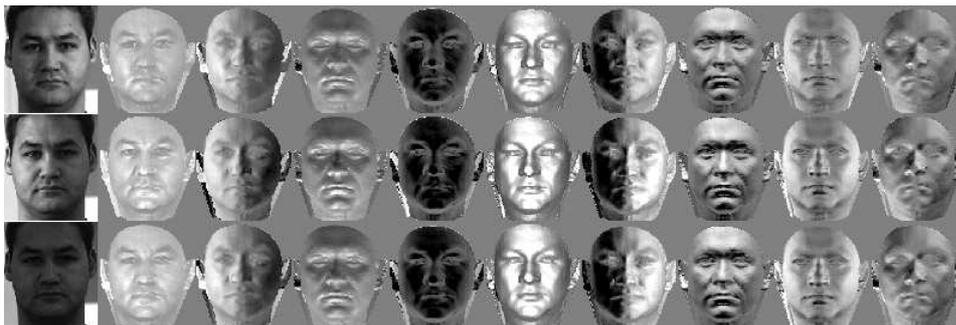


Fig. 4. Rendered basis images from training images taken under different illumination conditions. The first column shows the training images.

6.2 Recognition Experiments

In recognition experiments, we used the same set of 4488 images in CMU-PIE. We used only one image per subject to recover the 3D face model. We used the front and side galleries for training and all three pose galleries for testing. Notice that training images can have very different illumination conditions (unlike [23]). We performed recognition by using both all the 9 basis images and the first 4 basis images. We report our experimental results and comparison to [23] in Table 1. From the experimental results, we find that our method gives good recognition rates. When the poses of training and testing images are very different, our method is not as good as [23] because we only used a set of feature points to recover the shape information and the shape recovery is not accurate enough.

Table 2. Recognition results using various previous methods and our method on Yale Database B. Except for our method, the data were taken from [34]

Methods	Subset1,2	Subset3	Subset4
Eigenfaces	100	74.2	24.3
Linear Subspace	100	100	85
Cones-attached	100	100	91.4
9PL	100	100	97.2
Cones-cast	100	100	100
2D HIE	100	99.7	96.9
Our Method	100	100	97.2

We also performed experiments on the Yale Face Database B [3] and compared our recognition results with other methods for fixed frontal pose. The Yale Database contains images of 10 people with 9 poses and 64 illuminations per pose. We used 45*10 frontal face images for 10 subjects with each subject having 45 face images taken under different directional light sources. The data set is divided to 4 subsets following [15]. Table 2 compares our recognition rates with previous methods. As can be seen from Table 2, the results from our method are comparable with methods that require extensive training data per subject even though our method requires only one training image per pose. For fixed pose, the 2D HIE method in [34] performs almost as well as the our method, however the performance is very sensitive to accurate alignment of the faces.

6.3 Multiple Directional Illumination

As we mentioned, most of the previous methods are only applicable to single directional lighting. We study the performance of our method on images taken under multiple directional illumination sources to test our method under arbitrary illuminations. We synthesized images by combining face images in our data set and performed experiments on front and side galleries. For each subject, we randomly selected 2-6 images from the training data set and combined them together with random weights to simulate face images under multiple directional illumination sources(16 images per subject). We did experiments on the synthesized images both during training step and testing step. Table 3 shows the experimental results and we can see that our method also performed equally well under multiple sources of arbitrary direction.

Table 3. Experimental results of images under multiple directional illumination. "s" denotes images under single directional lighting and "m" denotes synthesized images under multiple illumination. "F" denotes the front gallery and "D" denotes the side gallery.

	Train:s; Test:s	Train:m; Test:s	Train:s; Test:m	Train:m; Test:m
Train:F; Test:F	98.2	98.3	97.8	98.1
Train:F; Test:D	92.4	92.0	91.5	92.2
Train:D; Test:F	94.4	93.6	94.2	94.8
Train:D; Test:D	96.7	95.9	96.3	96.1

7 Conclusions and Future Work

We have shown that by combining a morphable model and spherical harmonic lighting representation, we can recover both shape and texture information from one single image taken under arbitrary illumination conditions. Experimental results indicate that our method's recognition rates are comparable to other methods for pose variant images under single illumination. Moreover our method performs as well in the case of multiple illuminants, which is not handled by most previous methods. During the training phase, we only need one image per subject without illumination and pose limitations to recover the shape and texture information. Thus, the training set can be expanded easily with new subjects, which is desirable in a Face Recognition System.

In our experiments, we tested both images under single- and multiple- directional illuminations. At this time, there exist relatively few publicly available sets of images of faces under arbitrary illumination conditions, so we plan to continue validation of our method with a database with more types of light sources, e.g. area sources. There is human interactivity in the initialization of the model and the feature correspondences. We plan to integrate head pose estimation methods [32] for model initialization and optical flow algorithms for shape error estimation. In the face recognition phase, our method needs to search the whole pose space, we expect great speed-up with a pre-filter process (again using face pose estimation algorithms) to narrow the search space.

References

1. Y. Adini, Y. Moses, and S. Ullman. Face recognition: The problem of compensating for changes in illumination direction. *PAMI*, 19(7):721–732, July 1997.
2. R. Basri and D.W. Jacobs. Lambertian reflectance and linear subspaces. *PAMI*, 25(2):218–233, February 2003.
3. P.N. Belhumeur, J.P. Hespanha, and D.J. Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. *PAMI*, 19(7):711–720, 1997.
4. P.N. Belhumeur and D.J. Kriegman. What is the set of images of an object under all possible illumination conditions. *IJCV*, 28(3):245–260, July 1998.
5. J. Bigun, K. Choy, and H. Olsson. Evidence on skill differences of women and men concerning face recognition. *AVBPA*, pages 44–51, 2001.
6. V. Blanz and T. Vetter. A morphable model for the synthesis of 3d-faces. In *SIGGRAPH*, 1999.
7. V. Blanz and T. Vetter. Face recognition based on fitting a 3d morphable model. *PAMI*, 25(9):1063–1074, Sept. 2003.
8. R. Cappelli, D. Maio, and D. Maltoni. Subspace classification for face recognition. *Biometric Authentication*, pages 133–142, 2002.
9. K.I. Chang, K.W. Bowyer, and P.J. Flynn. Multi-modal 2d and 3d biometrics for face recognition. *AMFG*, pages 187–194, 2003.
10. R. Chellappa, C.L. Wilson, and S. Sirohey. Human and machine recognition of faces: A survey. *PIEEE*, 83(5):705–740, May 1995.
11. A.S. Georghiades, P.N. Belhumeur, and D.J. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *PAMI*, 23(6):643–660, June 2001.

12. P.W. Hallinan. A low-dimensional representation of human faces for arbitrary lighting conditions. *CVPR*, pages 995–999, 94.
13. R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.
14. B.W. Hwang, V. Blanz, T. Vetter, and S.W. Lee. Face reconstruction from a small number of feature points. In *ICPR*, pages Vol II: 842–845, 2000.
15. K.C. Lee, J. Ho, and D.J. Kriegman. Nine points of light: Acquiring subspaces for face recognition under variable lighting. In *CVPR*, pages I:519–526, 2001.
16. S.Z. Li and J.W. Lu. Face recognition using the nearest feature line method. *Neural Networks*, 10(2):439–443, 1999.
17. S.Z. Li, J. Yan, X. Hou, Z. Li, and H. Zhang. Learning low dimensional invariant signature of 3-d object under varying view and illumination from 2-d appearances. *ICCV*, pages 635–640, 2001.
18. C. Liu and H. Wechsler. A gabor feature classifier for face recognition. *ICCV*, pages 270–275, 2001.
19. X.M. Liu, T. Chen, and B.V.K.V. Kumar. Face authentication for multiple subjects using eigenflow. *Pattern Recognition, Special issue on Biometric*, Nov 2001.
20. X. Lu and A.K. Jain. Resampling for face recognition. *AVBPA03*, pages 869–877.
21. H. Murase and S.K. Nayar. Visual learning and recognition of 3-d objects from appearance. *IJCV*, 14(1):5–24, January 1995.
22. R. Ramamoorthi. Analytic pca construction for theoretical analysis of lighting variability in images of a lambertian object. *PAMI*, 24(10), Oct. 2002.
23. S. Romdhani, V. Blanz, and T. Vetter. Face identification by fitting a 3d morphable model using linear shape and texture error functions. In *ECCV*, 2002.
24. M Sadeghi, J. Kittler, A. Kostin, and K. Messer. A comparative study of automatic face verification algorithms on the banca database. *AVBPA*, pages 35–43, 2003.
25. A. Shashua. Illumination and view position in 3d visual recognition. In *NIPS*, 1991.
26. A. Shashua and Riklin Raviv. T. The quotient image: Class based re-rendering and recognition with varying illuminations. *PAMI*, 23(2):129–139, 2001.
27. T. Sim, S. Baker, and M. Bsat. The cmu pose, illumination, and expression (pie) database of human faces. In *AFGR*, pages 46–51, 2002.
28. M. Tistarelli, A. Lagorio, and E. Grosso. Understanding iconic image-based face biometrics. *Biometric Authentication*, pages 19–29, 2002.
29. M. Turk and A.P. Pentland. Eigenfaces for recognition. *CogNeuro*, 3(1):71–96, 1991.
30. H. Wang, S.Z. Li, Y. Wang, and W. Zhang. Illumination modeling and normalization for face recognition. In *Proc ICCV Workshop on Analysis and Modeling of Faces and Gestures*, 2003.
31. Y. Wang, T. Tan, and A.K. Jain. Combining face and iris biometrics for identity verification. *AVBPA*, pages 805–813, 2003.
32. Y. Wei, L. Fradet, and T. Tan. Head pose estimation using gabor eigenspace modeling. In *ICIP*, pages I: 281–284, 2002.
33. J. Yang, D. Zhang, and A.F. Frangi. Two-dimensional pca: A new approach to appearance-based face representation and recognition. *PAMI*, 26(1):131–137, 2004.
34. L. Zhang and D. Samaras. Face recognition under variable lighting using harmonic image exemplars. In *CVPR*, pages I: 19–25, 2003.
35. W. Zhao, R. Chellappa, A. Rosenfeld, and P.J. Phillips. Face recognition: A literature survey. In *UMD Technical Report CAR-TR948*, 2000.
36. W.Y. Zhao and R. Chellappa. Illumination-insensitive face recognition using symmetric shape-from-shading. In *CVPR*, pages I: 286–293, 2000.