Shape-from-shading driven 3D Morphable Models for Illumination Insensitive Face Recognition

Ankur Patel
ankur@cs.york.ac.uk
William A.P. Smith
wsmith@cs.york.ac.uk
Department of Computer Science,
The University of York

Abstract

In this paper we present a method for face shape and albedo estimation which uses a morphable model in conjunction with non-Lambertian shape-from-shading. We use surface normal and albedo estimates to construct a spherical harmonic basis which can be used generatively to model face appearance variation under arbitrarily complex illumination. This allows us to perform illumination insensitive face recognition given only a single gallery image. In contrast to other similar methods, our surface normal and albedo estimates are not constrained by a statistical model and are instead inferred from shading cues. We present recognition results on the Yale Face Database B.

1 Introduction

The most challenging formulation of the face recognition problem is to recognise a face, previously seen only once, under radically different pose or illumination conditions. In this paper we focus on extreme variations in illumination. Here, the appearance of the same subject can vary dramatically when the lighting direction changes. In fact, entirely different portions of the face may be visible under two different extreme illuminations since much of the face will be in shadow.

Where multiple images of a subject under varying illumination are available at the training stage, a number of appearance-based approaches have proven robust for illumination insensitive recognition [1, 6, 7]. The idea is to model image variability caused by changes in illumination using a linear subspace constructed for each subject. The basis set can be used in a generative manner to synthesise photorealistic images under arbitrary and possibly extreme lighting conditions. Impressive illumination insensitive face recognition performance can be achieved using these approaches. The drawback of these approaches is that they either require multiple training images (typically 7-9) or knowledge of the underlying shape and reflectance information (which may be recovered from the multiple training images). Further, appearance-based approaches cannot be used for pose invariant recognition as the constructed subspace is valid only for a single pose.

For the harder case, where only a single training image is available, the most promising methods use 3D face shape information. This has the obvious advantage that it is an intrinsic

BMVC 2009 doi:10.5244/C.23.22

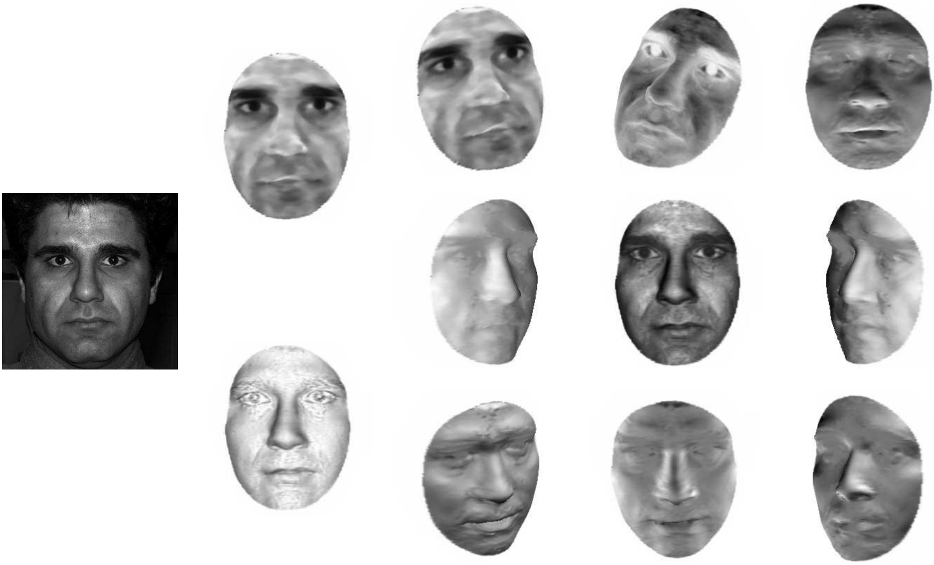

Figure 1: Given an input 2D image under frontal illumination (first column). The figure shows the estimated albedo map and bump map (second column) using the algorithm described in this paper. The last 3 columns show the spherical harmonic subspace derived from the estimated albedo and bump maps.

property of the face and hence is invariant to both illumination conditions and pose. The challenge is to recover accurate 3D facial shape information from a single image. The morphable model of Blanz and Vetter [2] captures variation in both 3D shape and texture and is fitted to a single image using an analysis-by-synthesis framework. The difficulty here is that fitting the model to an image requires the costly minimisation of an error functional, the solution of which suffers from model dominance. For recognition, the model must be fitted to both gallery and probe images and the parametric descriptors of the two compared. Zhang and Samaras [13] on the other hand, use the fit of a morphable model to each gallery image to obtain a spherical harmonic basis for that subject. This is also done in a model-based manner by learning a statistical model of spherical harmonic bases. The resulting basis can be used generatively to predict the appearance of a subject under arbitrarily complex illumination, although an assumption of Lambertian reflectance is made.

An interesting technique which does not fit into either of the classifications above uses local *gradient angle* features extracted from a single training image, which are subsequently matched against features from probe images [4]. The idea is to learn which such features are insensitive to variations in illumination using a large dataset of objects observed under varying illumination. However, as for the appearance-based approaches, this method cannot handle varying pose.

In this paper, we present a non-Lambertian face shape estimation algorithm and show how it can be used for illumination insensitive face recognition. Our method combines ideas from morphable models with those from classical shape-from-shading. The aim is to retain the robustness and flexibility of using a statistical model, with the fine surface detail and discriminating features conveyed by irradiance cues. Our algorithm uses a morphable model

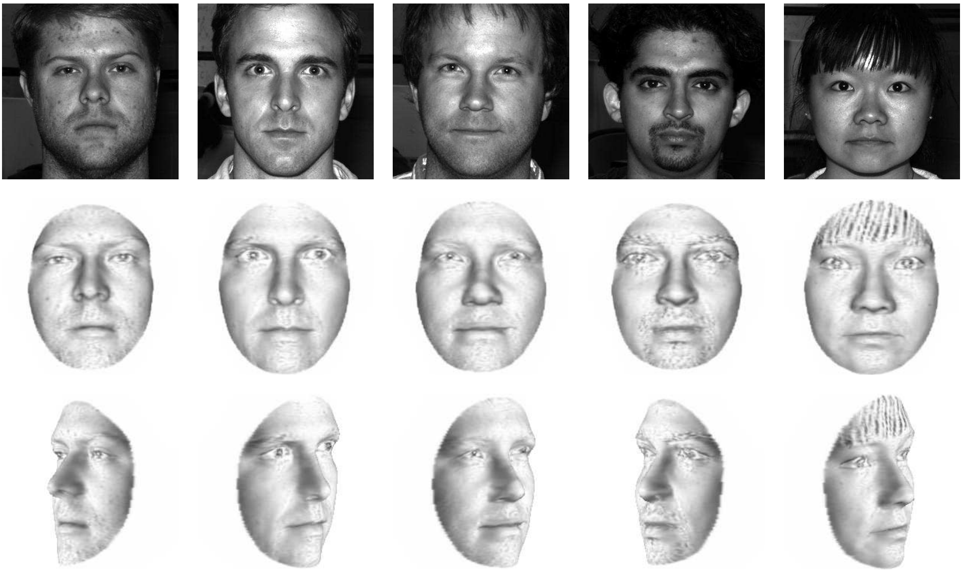

Figure 2: Shows the estimated bump maps rendered on the fitted morphable model meshes for 5 subjects in the Yale Face Database B.

to obtain a 3D face mesh and shape-from-shading to estimate a non-model-based surface normal map (bump map) and diffuse albedo map. The surface normal and diffuse albedo maps not are constrained by a statistical model and are therefore free to capture atypical, discriminating facial features. We perform shape recovery on the gallery images only and use the surface normals and albedo estimates to compute a spherical harmonic basis which is subsequently used for recognition. We present results on the Yale Face Database B [6] and compare our method with state-of-the-art published results.

2 3D Morphable Models as Shape Spaces

We use a morphable model to capture variations in the 3D shape of a human face. The morphable model is constructed according to Kendall's notion of a shape space [10].

Any face surface $\mathbf{x}$ can be represented as a linear combination of the average surface and the model eigenvectors P_i

$$\mathbf{x} = \bar{\mathbf{x}} + \sum_{i=1}^{m} b_i P_i = \bar{\mathbf{x}} + \mathbf{P}\mathbf{b}, \quad (1)$$

where $\mathbf{b} = (b_1, \ldots, b_m)^T$ is a vector of shape parameters and m is the number of face meshes in the training set. $\mathbf{P}$ is the matrix of the n most significant eigenvectors. We may choose to retain $n < m$ model dimensions, such that a certain percentage of the cumulative variance, described by the eigenvalues λ_i is captured.

The lengths of the parameter vectors, as measured by the square of the Mahalanobis

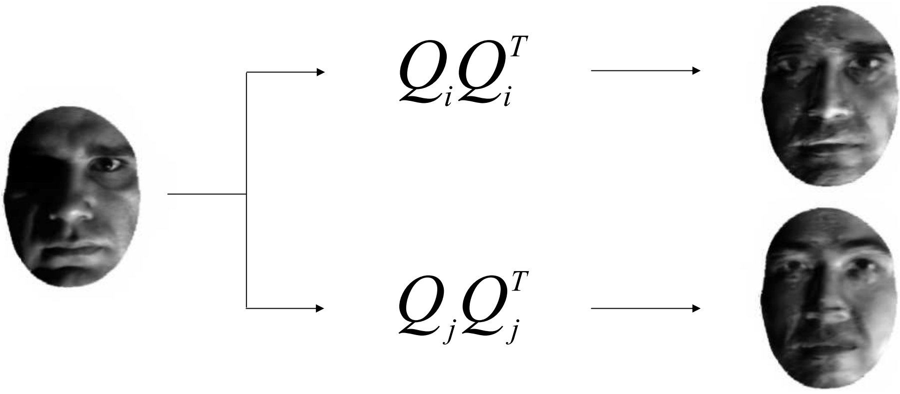

Figure 3: The figure explains the recognition principle. Given an input image (left column) of subject i, captured under lighting angles of $50° - 77°$. The right column shows the images projected onto the i^{th} spherical harmonic subspace (top) and the j^{th} spherical harmonic subspace (bottom). Note j corresponds to to a subject in the database disjoint from i

distance from the mean follow a chi-square distribution with n degrees of freedom [10]:

$$D_M^2(\mathbf{b}) = \sum_{i=1}^{n} \left(\frac{b_i}{\sqrt{\lambda_i}} \right)^2 \sim \chi_n^2. \tag{2}$$

An interesting observation is that the chi-squared distribution of parameter vector lengths implies that the parameter vectors lie approximately on the surface of hyperellipsoid in parameter space. This observation suggests sensible constraints to enforce on the parameter vector lengths [10].

3 Shape-from-Shading

The aim of computational shape-from-shading is to estimate 3D surface shape from single 2D intensity images. In order to recover surface orientation from image intensity measurements, the reflectance properties of the surface (human skin in our case) must be modelled.

The Blinn-Phong reflectance model [3] is a phenomenological attempt to describe surfaces which reflect light both specularly and diffusely. It comprises a Lambertian diffuse term and a specular term controlled by the shininess parameter:

$$g_{Phong}(\theta_i, \theta_h, \rho_d, \rho_s, \eta_s) = \rho_d \cos(\theta_i) + \rho_s \cos^{\eta_s}(\theta_h), \tag{3}$$

where ρ_d is the diffuse albedo, θ_i is the angle between the surface normal $\mathbf{n}$ and light source $\mathbf{s}$ vectors, ρ_s is the specular coefficient, η_s is the shininess parameter and θ_h is the angle between the surface normal $\mathbf{n}$ and the vector $\mathbf{h} = \frac{\mathbf{s}+\mathbf{v}}{\|\mathbf{s}+\mathbf{v}\|}$ which bisects the light source $\mathbf{s}$ and viewer $\mathbf{v}$ directions.

For an image in which the viewer and light source directions are fixed, the image irradiance equation (3) reduces to a function of the surface normal direction. For typical reflectance models, this equation does not have an unique minimum and there are likely to be an infinite set of normal directions all of which minimise the equation.

The Lambertian diffuse term in 3, provides a partial constraint on the direction of the surface normal, namely that the angle between the light source and the surface normal is given by:

$$\theta_i = \arccos(\mathbf{n}\cdot\mathbf{s}) = \arccos\left(\frac{I}{\rho_d}\right). \quad (4)$$

Geometrically, this means that the surface normal must lie on a right circular cone whose axis is the light source direction and whose half angle is θ_i. By constraining the surface normal to lie on the cone, the image irradiance equation is strictly satisfied, thereby ensuring that the information conveyed by the image is used to its fullest extent. Worthington and Hancock [12] show how to restore this constraint by rotating a surface normal to its closest on-cone position:

$$\acute{\mathbf{n}} = \Theta(\mathbf{a},\alpha)\,\tilde{\mathbf{n}}, \quad (5)$$

where, $\tilde{\mathbf{n}}$ is an off-cone surface normal and Θ is a rotation matrix which rotates a unit vector about axis $\mathbf{a}$ by an angle α. To restore a normal to the cone, we set $\mathbf{a} = \tilde{\mathbf{n}}\times\mathbf{s}$ and $\alpha = \theta_i - \arccos[\tilde{\mathbf{n}}\cdot\mathbf{s}]$. $\acute{\mathbf{n}}$ is the closest on-cone position that satisfies $\theta_i = \arccos\left(\frac{I}{\rho_d}\right)$.

We can rearrange 3 and reformulate 4 :

$$\theta_i = \arccos(\mathbf{n}\cdot\mathbf{s}) = \arccos\left(\frac{I-\rho_s\cos^{\eta_s}(\theta_h)}{\rho_d}\right). \quad (6)$$

If we use this angle in the computation of the rotation in (5), we strictly enforce the image irradiance constraint on a surface normal $\acute{\mathbf{n}}$ by finding the closest surface normal direction that satisfies $\theta_i = \arccos\left(\frac{I-\rho_s\cos^{\eta_s}(\theta_h)}{\rho_d}\right)$.

4 Shading Constraints for Morphable Models

A morphable model allows us to represent a novel face using a linear combination of an orthonormal basis. Shape-from-shading enables us to modify an estimated set of surface normals such that they strictly satisfy constraints implied by the reflectance properties of the surface. In this section we combine these two ideas and formulate a framework that enables us to estimate the 3D shape based on minimising the arc distance $d(\mathbf{e},\mathbf{f}) = \arccos(\mathbf{e}\cdot\mathbf{f})$, between two unit normals:

$$\mathbf{b}^* = \underset{D_M^2(\mathbf{b})\le D_{\max}^2}{\arg\min} \sum_p d\left(\mathbf{n}^p(\mathbf{b}), \Theta\mathbf{n}^p(\mathbf{b})\right), \quad (7)$$

where $\mathbf{n}^p(\mathbf{b})$ is the pth vertex normal obtained from the 3D Morphable model. Θ is the rotation matrix that strictly enforces satisfaction of the image irradiance equation on the pth vertex normal. $\mathbf{b}^*$ is the optimum shape parameter vector that minimises the expression. $D_{\max}^2$ is the maximum allowable parameter vector length, which controls the trade off between fitting quality and shape plausibility. Optimum performance occurs when $D_{\max}^2 \approx n$ [10]. Vertex normals are computed from the mesh which results from the shape parameters $\mathbf{b}$ using Max's [9] algorithm:

$$\mathbf{n}^p(\mathbf{b}) = ||\sum_f \mathbf{n}^f(\mathbf{b})\left|\mathbf{E}_f\times\mathbf{E}_{f+1}\right|, \quad (8)$$

where the summation is over all f faces incident to the vertex in question and $\mathbf{n}^f(\mathbf{b})$ is the face normal of the fth face. $\left|\mathbf{E}_f \times \mathbf{E}_{f+1}\right|$ provides weights based on areas of the adjacent triangles. $\|$ represents the normalisation step.

For each vertex in the mesh, we sample the image intensity by projecting the vertex to the image plane using an orthographic projection. Note that since vertices do not correspond directly to pixels, we use subpixel sampling to associate an intensity to each vertex. We denote the intensity associated with the pth vertex as $I(\hat{r}_p)$. It is these intensity values which provide a constraint on the surface normal, which is strictly satisfied by applying the rotation $\Theta\mathbf{n}^p(\mathbf{b})$. The rotated vertex normals strictly satisfy (6). These shape-from-shading normals possess two important qualities: 1. they will exactly recreate the input image 2. they are not constrained by the statistical model. The result of this is that they will capture fine surface detail.

The final ingredient in our method is to iteratively update the parameters of our reflectance model. Rearranging (6), we can obtain a per vertex estimate of the diffuse albedo. However, there is an additional constraint we may impose here. The diffuse albedo may not be greater than 1 (since a surface cannot reflect more light than was incident upon it). Hence

$$\rho_d(\hat{r}_p) = \min\left(1, \frac{I(\hat{r}_p) - \rho_s \cos^{\eta_s}(\theta_h^p(\mathbf{b}))}{\mathbf{n}^p(\mathbf{b}) \cdot \mathbf{s}}\right). \tag{9}$$

We use the estimated surface normal map (bump map) and albedo map to formulate our recognition principle.

5 Model Fitting Implementation

Our implementation estimates morphable model shape parameters using an optimisation based on shape-from-shading constraints. The input to our algorithm is a single intensity image, the light source direction and the viewer direction. We make a number of assumptions to simplify the fitting of the model. The first follows [5] and assumes that the specular coefficient (ρ_s) and roughness parameter η, are constant over the surface. We allow the diffuse albedo, ρ_d, to vary spatially over the face. Allowing the albedo and surface normal to vary arbitrarily renders the problem under-constrained. Therefore, in practice we enforce an additional regularisation constraint which requires the albedo to be piecewise smooth (by applying an edge sensitive smoothing over a small region).

Following Blanz and Vetter [2], we initialise our optimisation using a coarse initial fit of the model to a sparse set of k 2D annotations ($L_{2d} \in \mathbb{R}^{2k}$) on the subject's face ($k << p$) [10]. This initialisation also provides an estimate of the 3D pose (γ).

$$\begin{aligned}
(\mathbf{b}^{initial}, \gamma^*) &= \underset{D_M^2(\mathbf{b}) \le D_{\max}^2, \gamma}{\arg\min} \quad E(\mathbf{b}, \gamma). \\
E(\mathbf{b}, \gamma) &= \left\| L_{2d} - \hat{L_{2d}} \right\|. \\
\hat{L_{2d}} &= P_o T_r^{-1}(\bar{\mathbf{x}} + \mathbf{P}\mathbf{b}, \gamma).
\end{aligned} \tag{10}$$

Using the above shape estimate we determine the initial albedo map $\rho_d^{initial}$ using (9). We initialise the specular parameters as $\rho_s = 0.2$ and $\eta = 20$.

The objective function at iteration t of the optimisation is given in Algorithm 1. We minimise this objective function using Levenberg-Marquardt's method [8] to find optimal estimates of the shape parameter vector. Our convergence criterion is based on the total angular error between the model-based and shape-from-shading normals:

$$\sum_p \left[\arccos\left(\mathbf{n}^p(\mathbf{b}) \cdot \Theta\mathbf{n}^p(\mathbf{b})\right)\right]^2 < \varepsilon. \tag{11}$$

6 Face Recognition

In [1, 6, 7] it has been shown that a low-dimensional subspace can accurately capture the variation in images of a face resulting from arbitrarily complex variations in illumination. This provides a powerful representation for recognition, in which the identity associated with the subspace which lies closest to a query image is reported as the unknown identity. A variety of approaches of varying complexity have been proposed to build these subspaces. However, all require either the acquisition of a number of training images or knowledge of the underlying shape and reflectance properties of the face. In this section we show how the albedo and bump maps estimated from a single image (under frontal illumination) can be used to build such subspaces.

In this work we follow the approach of Basri and Jacobs [1] based on spherical harmonics in which the low-dimensional subspace is derived analytically from a model. They show that under any lighting conditions, at least 98% of the variability in the reflectance function is captured by the first 9 harmonic images. Their analysis therefore suggests that images of a convex Lambertian surface will lie close to a 9D subspace. This subspace can be derived exactly from the estimated albedo and bump maps without being dependent on the quantity or variability of a sample of training images.

Let ρ denote a vector of length M containing the albedo values across a face's surface, such that ρ_q is the albedo at point q. Similarly, the x, y and z components of the surface normals are stacked to form a further three vectors of length M: n_x, n_y and n_z, such that $n_{x,q}$ is the x component of the surface normal at point q. We define: $n_{x^2} = n_x. * n_x$ (where the operator $.*$ denotes the component-wise product of two vectors of the same length). Similarly for n_{y^2}, n_{z^2}, n_{xz}, n_{yz} and n_{xy}. The first nine harmonic images for a surface with known normals and albedo are given by:

$$\begin{aligned}
b_{00} &= \tfrac{1}{\sqrt{4\pi}}\rho, & b^e_{10} &= \sqrt{\tfrac{3}{4\pi}}\rho. * n_z, \\
b^o_{11} &= \sqrt{\tfrac{3}{4\pi}}\rho. * n_y, & b^e_{11} &= \sqrt{\tfrac{3}{4\pi}}\rho. * n_x, \\
b_{20} &= \tfrac{1}{2}\sqrt{\tfrac{3}{4\pi}}\rho. * (2n_{z^2} - n_{x^2} - n_{y^2}), \\
b^o_{21} &= 3\sqrt{\tfrac{5}{12\pi}}\rho. * n_{yz}, & b^e_{21} &= 3\sqrt{\tfrac{5}{12\pi}}\rho. * n_{xz}, \\
b^o_{22} &= 3\sqrt{\tfrac{5}{12\pi}}\rho. * n_{xy}, & b^e_{22} &= \tfrac{3}{2}\sqrt{\tfrac{5}{12\pi}}\rho. * (n_{x^2} - n_{y^2}).
\end{aligned} \tag{12}$$

It is clear that these harmonic images may be derived from precisely the information recovered by our algorithm. Once again, we form a matrix B containing the basis images as columns (this time of dimension $M \times 9$). However, this basis is not orthonormal. Using a QR decomposition, we find the $M \times 9$ orthonormal basis Q and 9×9 matrix R, such that $QR = B$. Given a vector of sampled image intensities $I(\hat{r}_p)$, we may now compute the distance to the subspace using: $\|QQ^T I(\hat{r}_p) - I(\hat{r}_p)\|$ and perform recognition as described above.

Algorithm 1: Objective Function for Fitting a 3D Morphable Model using Shape-from Shading

Input: Light source direction $\mathbf{s}$, Viewer direction $\mathbf{v}$, input image I and shape parameters $\mathbf{b}$

Output: Albedo map ρ_d and bump map $\mathbf{N}$

1 Constrain the shape parameter vector length (7);

2 Obtain the p vertex normals $\mathbf{n}^p(\mathbf{b})$ using (8) ;

3 Project vertices to image plane using orthographic projection to obtain vertex intensity estimates ;

4 Apply regularisation constraint on the diffuse albedo: $\rho_d^{(t-1)} = f(\rho_d^{(t-1)})$;

5 Update vertex normals according to sampled image intensities giving bump map $\mathbf{N}^{(t)} = \Theta \mathbf{n}^p(\mathbf{b})$;

6 Update the albedo map (9): $\rho_d^{(t)}(\hat{r}_p) = \min\left(1, \frac{I(\hat{r}_p) - \rho_s \cos^{\eta_s}(\theta_h^p(\mathbf{b}))}{\mathbf{n}^p(\mathbf{b}) \cdot \mathbf{s}}\right)$;

7 Experimental Results

In this section we present recognition results using the algorithm proposed in this paper. The 3D morphable model (Section 2) was built using 100 scans obtained from a Cyberware 3030PS laser range scanner. We retain 99 modes for our experiment.

We provide face recognition results on the data obtained from the Yale Face Database B [6], which contains images of 10 individuals (disjoint from the morphable model training data) under 45 different illumination conditions. We group the lighting variation into 4 subsets of differing extremity (see [6] for details). The images have been cropped using the metadata provided along with the dataset. All images are resized to 305x305 pixels.

For our recognition experiment we use a single training image with frontal lighting (azimuth and elevation angles equal to zero degrees). Although we could choose any illumination for the training image, we use frontal illumination for two reasons. Firstly, this configuration ensures that none of the face is in shadow. Non-frontal lighting would result in a degradation in the accuracy of the estimated albedo and bump maps and subsequent recognition performance. Secondly, this corresponds to a useful and realistic scenario in which the training image is captured using a standard camera with flash. We apply our algorithm (Section 5) to each training image and use the estimated albedo and bump maps to perform recognition (as described in Section 6). Figure 1 shows the estimated albedo map, bump map and spherical harmonic basis for a subject in the database. In Figure 2, we show images in a novel pose rendered with the estimated bump maps. It is clear that in all cases the shape recovery process is stable (even with large variations in albedo caused by facial hair) and that the bump maps successfully capture discriminating facial shape details. Figure 3 explains the recognition principle. Processing a gallery image requires implementation of our algorithm which takes takes 2 to 2.5 minutes to run on a 1.78 GHz AMD Athlon processor. Processing the probe images simply requires sampling of image intensities under an orthographic projection using the shape and pose parameters estimated for the gallery image plus the spherical harmonic reconstruction. Both of these steps can be implemented very efficiently. In contrast, the method of Blanz and Vetter [2] requires 4.5 minutes per gallery *and* probe image as the morphable model must be fitted from scratch in both cases. In addition, our method does not require any prior knowledge or estimation of the lighting conditions in a probe image (they can be arbitrarily complex), whereas Blanz and Vetter must explicitly estimate the lighting conditions when they fit to a probe image.

Table 1 shows the recognition results for our method and compares them with other methods. Extending our method to non-frontal pose is particularly straightforward, simply requiring the rigid alignment of the 3D gallery model to a non-frontal probe image. Also, since the estimated bump map and albedo map are not constrained by a model, they pick up discriminating features which are likely to be important for recognition in larger databases.

8 Conclusions

We have shown how ideas from shape-from-shading and morphable models can be combined in a robust face shape estimation framework. By using shading cues, we can obtain surface normal and albedo estimates that are not constrained by a statistical model and are therefore free to capture atypical and discriminating surface features and markings. We have shown how this data can be used to construct a spherical harmonic basis that can be subsequently used to perform illumination insensitive face recognition from a single gallery image. In future work we intend to experiment with larger datasets and variations in pose between probe and gallery images. We also aim to extend our face shape estimation algorithm, such that reconstruction can take place from images containing arbitrarily complex illumination.

References

[1] R. Basri and D. W. Jacobs. Lambertian reflectance and linear subspaces. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25(2):218–233, 2003.

[2] V. Blanz and T. Vetter. Face recognition based on fitting a 3D morphable model. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25(9):1063–1074, 2003.

[3] J. F. Blinn. Models of light reflection for computer synthesized pictures. *SIGGRAPH Comput. Graph.*, 11:192–198, 1977.

[4] H. Chen, P. Belhumeur, and D.Jacobs. In search of illumination invariants. In *Proc. CVPR*, pages 1–8, 2000.

[5] A. Georghiades. Recovering 3–d shape and reflectance from a small number of photographs. In *Eurographics Symposium on Rendering*, pages 230–240, 2003.

[6] 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. *IEEE Trans. Pattern Anal. Mach. Intell.*, 23(6):643–660, 2001.

[7] K. Lee, J. Ho, and D. Kriegman. Acquiring linear subspaces for face recognition under variable lighting. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(5):1–15, 2005.

[8] M. I. A. Lourakis. levmar: Levenberg-Marquardt nonlinear least squares algorithms in C/C++. http://www.ics.forth.gr/~lourakis/levmar/.

[9] Nelson Max. Weights for computing vertex normals from facet normals. *Journal of graphics tools*, 4(2):1–6, 1999.

Comparison of Recognition Methods					
Method	Number of Training Images	Error Rate(%) vs. Illum			
		Subset 1&2	Subset 3	Subset 4	Total
Correlation	6-7	0.0	23.3	73.6	29.1
Eigenfaces	6-7	0.0	25.8	75.7	30.4
Eigenfaces w/o 1st 3	6-7	0.0	19.2	66.4	25.8
3-D Linear subspace	6-7	0.0	0.0	15	4.6
Cones-attached	6-7	0.0	0.0	8.6	2.7
Harmonic Subspace-attached (no cast shadow)	6-7	0.0	0.0	3.571	1.1
Harmonic Subspace-attached (with cast shadow)	6-7	0.0	0.0	2.7	0.85
Cones-cast	6-7	0.0	0.0	0.0	0.0
5PL	5	0.0	0.0	0.0	0.0
9PL	9	0.0	0.0	0.0	0.0
Gradient Angle	1	0.0	0.0	1.4	0.44
Harmonic Exemplars *	1	0.0	0.3	3.1	1.0
Zhang & Samaras *	1	0.0	0.0	3.1	0.97
Our Method	1	0.0	0.0	5	1.59

Table 1: Recognition results on the Yale Face Database B. Except our method, all the remaining data was taken from [ ]. All the methods (except the ones with *) report results on 440 test images. The methods with * report results on 400 test images, with scores averaged over 10 experimental runs.

[10] A. Patel and W. A. P. Smith. 3D morphable face models revisited. In *Proc. CVPR*, pages 1327–1334, 2009.

[11] S. Romdhani, J. Ho, T. Vetter, and D. J. Kriegman. Face recognition using 3–d models: Pose and illumination. *Proc. of the IEEE*, 94(11):1977–1999, 2006.

[12] P. L. Worthington and E. R. Hancock. New constraints on data–closeness and needle map consistency for shape–from–shading. *IEEE Trans. Pattern Anal. Mach. Intell.*, 21 (12):1250–1267, 1999.

[13] L. Zhang and D. Samaras. Face recognition from a single training image under arbitrary unknown lighting using spherical harmonics. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(3):351–363, 2006.