

# Molding Face Shapes by Example

Ira Kemelmacher and Ronen Basri\*

Dept. of Computer Science and Applied Math.,  
The Weizmann Institute of Science,  
Rehovot 76100, Israel

**Abstract.** Human faces are remarkably similar in global properties, including size, aspect ratios, and locations of main features, but can vary considerably in details across individuals, gender, race, or due to facial expression. We propose a novel method for 3D shape recovery of a face from a *single* image using a *single* 3D reference model of a different person's face. The method uses the input image as a guide to mold the reference model to reach a desired reconstruction. Assuming Lambertian reflectance and rough alignment of the input image and reference model, we seek shape, albedo, and lighting that best fit the image while preserving the rough structure of the model. We demonstrate our method by providing accurate reconstructions of novel faces overcoming significant differences in shape due to gender, race, and facial expressions.

## 1 Introduction

The 3-dimensional shape of a face and its reflectance properties contain important information that can be used for recognition and for predicting appearance under novel viewing conditions. Recovering this information from a single image is difficult, since shape from shading algorithms generally require knowledge of the lighting conditions and the reflectance properties of the face [1, 2, 3, 4] (see some attempts to relax these assumptions in [5, 6, 7]). People, in contrast, seem to skillfully recognize faces from novel images overcoming significant viewpoint and lighting variations. This ability is often attributed to familiarity with faces as a class (e.g., [8]).

To address this difficulty, various algorithms use class information to restrict the set of allowable reconstructions. One approach attempts to exploit the symmetry of faces [9, 10]. The advantage of using symmetry is that reconstruction can rely on a mere single image without the need for additional examples of face models. The disadvantage is that point-wise correspondence between the two symmetric portions must be established, and this task is generally difficult. Another approach is to learn the set of allowable reconstructions from a large number of faces in a database. This can be achieved by embedding all 3D faces

---

\* Research was supported in part by the Israel Science Foundation grant number 266/02 and by the European Commission Project IST-2002-506766 Aim Shape. The vision group at the Weizmann Inst. is supported in part by the Moross Laboratory for Vision Research and Robotics.

in a linear space [11, 12, 13, 14, 15] (see also [16] where this approach is combined with symmetry) or by using a training set to determine a density function for faces [17, 18]. These methods can achieve accurate reconstructions, but they require a large number of face models as well as point-wise correspondence between all the models. Finally, [19] proposed a method for rendering faces in novel views assuming that different faces share the exact same shape while differ only in albedo.

In a global sense, different faces indeed are highly similar. Faces of different individuals share the same main features (eyes, nose, mouth) in roughly the same locations, and their sizes and aspect ratios do not vary much. However, locally, face shapes can vary considerably across individuals, gender, race, or as a result of facial expressions. Face recognition methods use this global similarity of faces, e.g., to estimate the pose of novel faces, for example by aligning a face image to a generic face model. In this paper we will demonstrate how this global similarity can be exploited to obtain a detailed shape reconstruction of novel faces.

Below we introduce a novel method for shape recovery of a face from a *single image* that uses only a *single reference 3D face model* of a different person in the training set. Intuitively, our method uses the input image as a guide to mold the reference model to reach a desired reconstruction. Specifically, the method modifies the shape and albedo of the model face to fit the image. Since in general selecting shape and albedo to fit an image is an ill posed problem, we will restrict the method to produce reconstructions that preserve the rough shape and albedo of the reference model.

Our method assumes Lambertian reflectance, light sources at infinity, and rough alignment between the input image and the reference model. It allows for multiple unknown light sources and attached shadows by using a spherical harmonic approximation to model reflectance (following [20, 21]). We cast the problem as an image irradiance equation [2] with unknown lighting, albedo, and surface normals. We then use the reference model to estimate lighting and provide initial estimate of albedo. We further introduce regularization terms to seek solutions that preserve the rough shape and albedo of the reference model. These terms will smooth the difference in shape and albedo between the reference model and the sought face. We show experiments demonstrating that the method can achieve accurate reconstructions of novel faces overcoming significant differences in shape due to gender, race, and facial expressions.

Although this paper emphasizes the use of a single model of a face to reconstruct another face, we note that this method can supplement methods that make use of multiple models in a database. In particular, we may select to mold the model from the database that best fits the image. Alternatively, we may choose the best fit model from a linear subspace spanned by the database, or we may choose a model based on probabilistic criteria. In all cases our method will try to improve the reconstruction by relying on the selected model.

The paper is divided as follows. Section 2 defines the optimization function. Section 3 describes the reconstruction algorithm. Experimental results are shown in Sect. 4.

## 2 Problem Statement

Consider an image  $E(x, y)$  of a face defined on a compact domain  $\Omega \subset \mathfrak{R}^2$ , whose corresponding surface is given by  $z(x, y)$ . The surface normal at every point is denoted  $\mathbf{n}(x, y)$  (boldface is used to denote vectors) with

$$\mathbf{n}(x, y) = \frac{1}{\sqrt{p^2 + q^2 + 1}}(p, q, -1)^T, \quad (1)$$

where  $p(x, y) = \partial z / \partial x$  and  $q(x, y) = \partial z / \partial y$ . We assume that the face is Lambertian with albedo  $\rho(x, y)$  and ignore the effect of cast shadows and inter-reflections. Under these assumptions, for an object illuminated by an arbitrary configuration of light sources at infinity, it has been shown [20, 21] that reflectance can be expressed in terms of spherical harmonics as

$$R(\mathbf{n}; \rho, \mathbf{l}) \approx \rho \sum_{i=0}^{K-1} l_i Y_i(\mathbf{n}), \quad (2)$$

where  $\mathbf{l} = (l_0, \dots, l_{K-1})$  denote the harmonic coefficients of lighting and  $Y_i(\mathbf{n})$  ( $0 \leq i < K - 1$ ) include the spherical harmonic functions evaluated at the surface normal. Because the reflectance of Lambertian objects under arbitrary lighting is very smooth this approximation is highly accurate already when a low order harmonic approximation is used. Specifically, a second order harmonic approximation (including nine harmonic functions) captures on average at least 99.2% of the energy in an image. A first order approximation (including four harmonic functions) can also be used with somewhat less accuracy. It has been shown analytically that a first order harmonic approximation captures at least 87.5% of the energy in an image, while in practice, owing to the fact that only normals with  $n_z \geq 0$  are observed, the accuracy seems to approach 95% [22]. Below we will model reflectance using a first order harmonic approximation and write this in vector notation as

$$R(\mathbf{n}; \rho, \mathbf{l}) \approx \rho \mathbf{l}^T \mathbf{Y}(\mathbf{n}), \quad (3)$$

with  $\mathbf{Y}(\mathbf{n}) = (1, n_x, n_y, n_z)^T$  and  $n_x, n_y, n_z$  are the components of  $\mathbf{n}^1$ .

The image irradiance equation is then given by

$$E(x, y) = R(\mathbf{n}; \rho, \mathbf{l}). \quad (4)$$

In general, when  $\rho$  and  $\mathbf{l}$  are provided this equation can be solved using shape from shading algorithms (e.g., [2, 3, 23, 24]), so we will need a method to estimate  $\rho$  and  $l$ .

To supply the missing information we will be assisted by a reference model of a face of a different individual. Let  $z_{\text{ref}}(x, y)$  denote the surface of the reference

<sup>1</sup> Formally, we should set  $\mathbf{Y} = (1/\sqrt{4\pi}, \sqrt{3/(4\pi)}n_x, \sqrt{3/(4\pi)}n_y, \sqrt{3/(4\pi)}n_z)$ . For convenience we omit these constant factors and rescale the lighting coefficients to include these factors.

face with  $\mathbf{n}_{\text{ref}}(x, y)$  denoting the normal to the surface, and  $\rho_{\text{ref}}(x, y)$  denote its albedo. We will use this information to determine the lighting and provide initial guess for the sought albedo.

Finally, to regularize the problem we will define the difference shape as

$$d_z(x, y) = z(x, y) - z_{\text{ref}}(x, y) \quad (5)$$

and the difference albedo as

$$d_\rho(x, y) = \rho(x, y) - \rho_{\text{ref}}(x, y) \quad (6)$$

and require that these differences will be smooth. We are now ready to define our optimization function:

$$\min_{\mathbf{l}, \rho, z} \int_{\Omega} ((E - \rho \mathbf{l}^T \mathbf{Y}(\mathbf{n}))^2 + \lambda_1 \Delta g(d_z) + \lambda_2 \Delta g(d_\rho)) dx dy. \quad (7)$$

$\Delta g(\cdot)$  denotes the Laplacian of a Gaussian function, and  $\lambda_1$  and  $\lambda_2$  are positive constants. Below we will refer to the first term in this integral as the “data term” and the other two terms as the “regularization terms”. Note that we chose to regularize  $d_z$  and  $d_\rho$  rather than  $z$  and  $\rho$  in order to preserve the discontinuities in  $z_{\text{ref}}$  and  $\rho_{\text{ref}}$ .

### 3 Surface Reconstruction

Evidently, without regularization the optimization functional (7) is ill-posed. Specifically, for every choice of depth  $z(x, y)$  and lighting  $\mathbf{l}$  it is possible to prescribe albedo  $\rho(x, y)$  to make the first term vanish. With regularization and appropriate boundary conditions the problem becomes well-posed.

We approach this optimization by solving for lighting, depth, and albedo separately. First, we recover the lighting coefficients  $\mathbf{l}$  by finding the best coefficients that fit the reference model to the image. This is analogous to solving for pose by matching the features of a model face to the features extracted from an image of a different face. Next we solve for depth  $z(x, y)$  using the recovered lighting coefficients and the albedo of the reference model. This in fact is the usual shape from shading problem. Finally, we use the lighting and the recovered depth to estimate the albedo  $\rho(x, y)$ . This procedure can be repeated iteratively, although in our experiments one iteration seemed to suffice. These three steps are described in detail in the next three subsections.

The use of the albedo of the reference model may seem restrictive since different people may vary significantly in skin color. Nevertheless, it can be readily verified that linearly transforming the albedo (i.e.,  $\alpha\rho(x, y) + \beta$ , with scalar constants  $\alpha$  and  $\beta$ ) can be compensated for by scaling appropriately the light intensity and changing the ambient term  $l_0$ . Our albedo recovery, consequently, will be subject to this ambiguity. It is important to note that to make sure that marks on the reference face would not influence much the reconstruction we first smooth the albedo of the reference model by a Gaussian.

### 3.1 Lighting Recovery

In the first step we attempt to recover the lighting coefficients by fitting the reference model to the image. To this end, we substitute in (7)  $\rho \rightarrow \rho_{\text{ref}}$  and  $z \rightarrow z_{\text{ref}}$  (and consequently  $\mathbf{n} \rightarrow \mathbf{n}_{\text{ref}}$ ). At this stage both regularization terms vanish, and only the data term remains:


$$\min_{\mathbf{I}} \int \int_{\Omega} (E - \rho_{\text{ref}} \mathbf{I}^T \mathbf{Y}(\mathbf{n}_{\text{ref}}))^2 dx dy. \quad (8)$$

Substituting for  $\mathbf{Y}$  and discretizing the integral we obtain

$$\min_{\mathbf{I}} \sum_{(x,y) \in \Omega} \left( E(x,y) - \rho_{\text{ref}}(x,y)(l_0 + \tilde{\mathbf{I}}^T \mathbf{n}_{\text{ref}}(x,y)) \right)^2, \quad (9)$$

where  $\tilde{\mathbf{I}} = (l_1, l_2, l_3)^T$ . This is a highly over-constrained linear least square optimization with only four unknowns (the components of  $\mathbf{I}$ ) and can be solved simply using the pseudo-inverse.

The lighting coefficients recovered with this procedure will be used subsequently to recover depth. To examine whether the coefficients recovered indeed are close to the true lighting coefficients we have run the following experiment. Using a database of 56 3D faces from the USF database [26] we recovered the lighting from images of each of these models by comparing the image to all the other 3D models in the database. We calculated for each such pair the angle between the true lighting and the recovered one; this represents the error in lighting recovery. The result of the experiment is shown in Fig. 1. We observe that the mean angle is  $11.3^\circ$  with standard deviation of  $6.2^\circ$ . As our experiments demonstrate (Sec. 4), this error is sufficiently small allowing accurate reconstructions.



**Fig. 1.** Accuracy of the lighting recovered. We plot a histogram of the angle (in degrees) between the true lighting coefficients and the recovered coefficients using reference models of different individuals. The distribution was calculated over 56 face shapes.

### 3.2 Depth Recovery

At this stage we have obtained an estimate for  $\mathbf{l}$ . We continue using  $\rho_{\text{ref}}$  for the albedo and turn to recovering  $z(x, y)$ . As we mentioned above,  $z$  can be recovered by solving a shape from shading problem, since the reflectance function is completely determined by the lighting coefficients and the albedo. Below we will further exploit the resemblance of the sought surface to the reference face to linearize the problem.

We first handle the data term. Denote by  $N(x, y) = \sqrt{p^2 + q^2 + 1}$ , we will assume that  $N(x, y) \approx N_{\text{ref}}(x, y)$ . The data term in fact minimizes the difference between the two sides of the following equation system

$$E = \rho_{\text{ref}} \left( l_0 + \frac{1}{N_{\text{ref}}} \tilde{\mathbf{l}}^T(p, q, -1)^T \right), \tag{10}$$

with  $p$  and  $q$  as unknowns. With additional manipulation this becomes

$$E - \rho_{\text{ref}} \left( l_0 - \frac{1}{N_{\text{ref}}} l_3 \right) = \frac{\rho_{\text{ref}}}{N_{\text{ref}}} (l_1 p + l_2 q). \tag{11}$$

In discretizing this equation system we will use  $z(x, y)$  as our unknowns, and replace  $p$  and  $q$  by the forward differences:

$$\begin{aligned} p &= z(x + 1, y) - z(x, y) \\ q &= z(x, y + 1) - z(x, y), \end{aligned} \tag{12}$$

obtaining

$$\begin{aligned} E - \rho_{\text{ref}} \left( l_0 - \frac{1}{N_{\text{ref}}} l_3 \right) = \\ \frac{\rho_{\text{ref}}}{N_{\text{ref}}} (l_1 (z(x + 1, y) - z(x, y)) + l_2 (z(x, y + 1) - z(x, y))). \end{aligned} \tag{13}$$

The data term thus provides one equation for every unknown. Note that by solving for  $z(x, y)$  we in fact enforce integrability.

Next we treat the regularization term  $\lambda_1 \Delta g(d_z)$  (the second regularization term vanishes at this stage). We implement this term as the difference between  $d_z(x, y)$  and the average of  $d_z$  around  $(x, y)$  obtained by applying a Gaussian function to  $d_z$  (denoted  $g(d_z)$ ). Consequently, this term minimizes the difference between the two sides of the following equation system

$$\lambda_1 (z(x, y) - g(z)) = \lambda_1 (z_{\text{ref}}(x, y) - g(z_{\text{ref}})). \tag{14}$$

It should be noted that to avoid degeneracies the input face must be lit by non-ambient light, since under ambient light intensities are independent of surface orientation. The assumption we used, that  $N(x, y) \approx N_{\text{ref}}(x, y)$  further requires that there will be light coming from directions other than the camera direction. If a face is lit from the camera direction (e.g., flash photography) then  $l_1 = l_2 = 0$

and the right-hand side of (11) vanishes. This degeneracy can be addressed by solving instead a usual nonlinear shape from shading algorithms (e.g., [3, 23, 24]).

Combining these two sets of equations we obtain a linear set of equations with two linear equations for every unknown. This system of equations is still rank deficient, and we need to add boundary conditions. We can use Dirichlet boundary conditions, but these will require us to know the depth values along the boundary of the face. We could use the depth values of the reference model, but these may be incompatible with the sought solution. Alternatively, we can constrain the derivatives of  $z$  along the boundaries using Neumann boundary conditions. One possibility is to assign  $p$  and  $q$  along the boundaries to match the corresponding derivatives of the reference model  $p_{\text{ref}}$  and  $q_{\text{ref}}$  so that the surface orientation of the reconstructed face along the boundaries will coincide with the surface orientation of the reference face. A less restrictive assumption is to assume that the surface is planar along the boundaries, i.e., that the partial derivatives of  $p$  and  $q$  in the direction orthogonal to the boundary  $\partial\Omega$  vanish. (Note that this does not imply that the entire boundaries are planar.) This assumption will be roughly satisfied if the boundaries are placed in slowly changing parts of the face. It will not be satisfied for example when the boundaries are placed along the eyebrows, where the surface orientation changes rapidly. We use this type of Neumann boundary conditions in our experiments.

Finally, since all the equations we use for the data term, the regularization term, and the boundary conditions involve only partial derivatives of  $z$ , while  $z$  itself is absent from these equations, the solution can be obtained only up to an additive factor. We will rectify this by arbitrarily setting one point to  $z(x_0, y_0) = z_0$ .

### 3.3 Estimating Albedo

Once both the lighting and depths are recovered, we may turn to estimating the albedo. Using the data term the albedo is given by

$$\rho(x, y) = \frac{E(x, y)}{l_0 + \tilde{\mathbf{I}}^T \mathbf{n}(x, y)}. \quad (15)$$

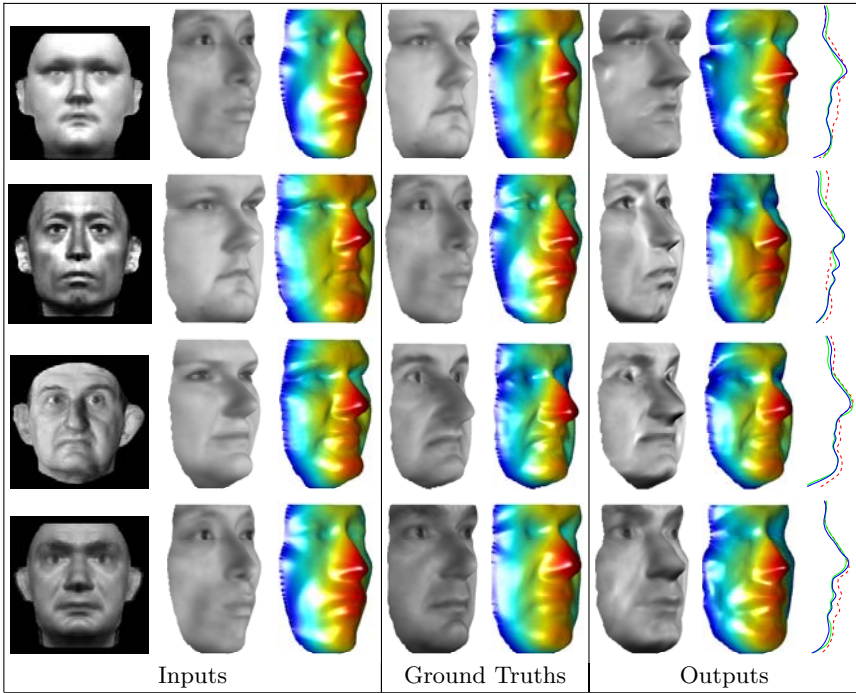
The first regularization term is independent of  $\rho$ , and so it can be ignored, and the second term optimizes the following equations

$$\lambda_2 \Delta g(\rho) = \lambda_2 \Delta g(\rho_{\text{ref}}). \quad (16)$$

Again these provide a linear set of equations, in which the first set determines the albedo values, and the second set smoothes these values. Boundary conditions are placed by simply terminating the smoothing process at the boundaries.

## 4 Experiments

To test our method we performed several sets of experiments. For reference models we used the first set of the USF face database, which contains depth and texture

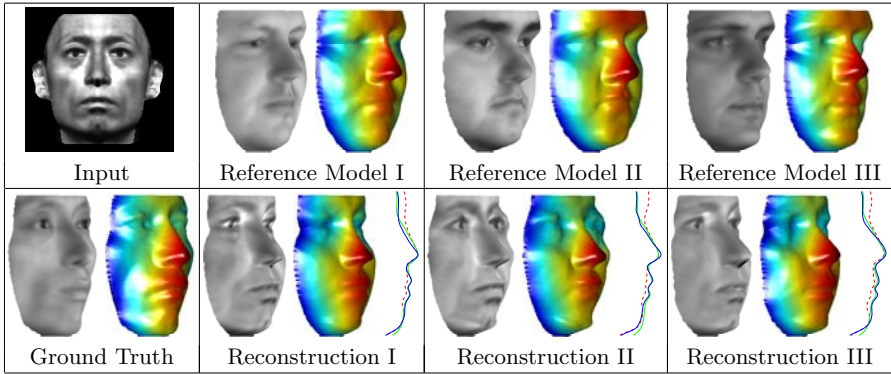


**Fig. 2.** Reconstruction from synthetic images. From left to right: Images rendered from the USF database, reference models (the surfaces are colored from blue to red according to  $z(x,y)$ ), and albedo painted on the model. These were used as inputs to our method. Ground truth shapes and albedos. The output obtained includes the recovered 3D shape and the recovered albedo painted on the output shape. Finally a profile curve of the recovered shape (blue) overlaid on the profile curve of the ground truth shape (green) and the profile curve of the reference model (red, dashed).


maps of 56 real faces (male and female adult faces with a mixture of race and age) obtained with a laser scanner [26]. The texture maps provided in USF database are not identical to the real albedos of the faces, since they contain noticeable effects of the lighting conditions. To reduce these effects we averaged each texture map with its mirror image, and used the result as albedos of the reference models.

In all experiments we attempted to recover the shape of frontal facing faces. The following parameters were used throughout all our experiments. The reference albedo was kept in the range between 0 and 255. Both  $\lambda_1$  and  $\lambda_2$  were set to 110. We smoothed the reference albedo by a 2-D Gaussian with  $\sigma_x = 3$  and  $\sigma_y = 4$ . The same smoothing parameters were used for the two regularization terms. Finally, to align the images with the reference models we marked five corresponding points on the image and the reference model, two at the centers of the eyes, one on the tip of the nose, one in the center of the mouth and one in the bottom of the chin (Fig. 4, right column). We then used these correspondences to determine a 2D rotation, translation, and scale to fit the image to the





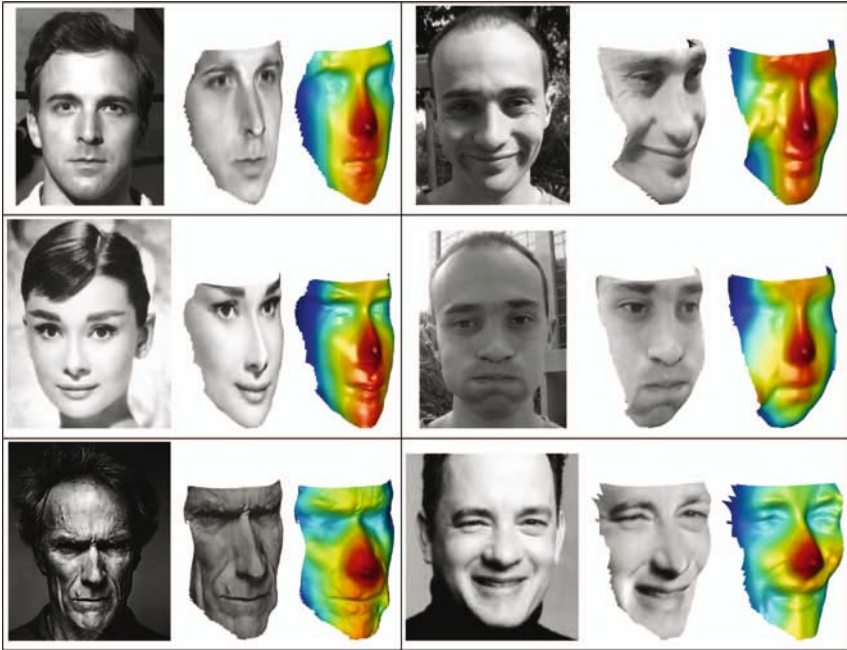
**Fig. 3.** Reconstructions of the same face using several different reference models. The first row contains the input image (left column) and three different reference models used as input. The second row contains the ground truth shape (left column) and the three reconstructions obtained using each of the reference models. An overlay of profiles is shown on the right of the each reconstruction (the recovered profile in blue, the ground truth profile in green, and the reference profile in dashed red).



**Fig. 4.** Left column: The model used for reference in the experiments with real images (Fig. 5). Right column: Five points used for alignment (two at the centers of the eyes, one on the tip of the nose, one in the center of the mouth and one in the bottom of the chin).

reference model. After alignment all the images contained  $150 \times 200$  pixels. To recover depth (Eqs. (13) and (14)) we directly solved a system of linear equations. Our non-optimized MATLAB implementation of the algorithm takes only 30 seconds on a Pentium IV PC.

The first set of experiments contain controlled experiments in which we artificially rendered faces from the USF database and then used our algorithm to recover their shapes and albedos from the rendered images. These experiments allow us to show comparisons of our reconstructions to the ground truth shapes. To produce an image we illuminated a model by 2-3 point sources from directions  $\mathbf{l}_i$  and with intensity  $L_i$ . The intensities reflected by the surface due to this light are given by  $I = \sum_{i=1}^n \rho L_i \max(\cos(\mathbf{n}^T \mathbf{l}_i), 0)$ . Fig. 2 shows several images obtained this way. For each image we selected a reference model of a different individual and used the image and the reference model to recover the depth and albedo



**Fig. 5.** Six experiments with real images. In each experiment, the input image and the reconstruction results are presented. Images were obtained from the YaleB database (top left), cropped from [12] (middle left), <http://www.swirc.com> (bottom left) and <http://crazy4cinema.com/Actor/hanks.html> (bottom right). The rest of the images were photographed by us.

of the rendered face. For comparison we show the reconstructed shapes and the laser scanned shapes. We show both the reconstructed and the scanned shapes in two ways, with albedo painted on the shape and in a colored representation with the color representing the depth values. The latter representation better displays the details of the shape independent of the variations in albedo. We further plot an overlay of the profile curves of the reconstructed shape (in blue), the ground truth model (green), and the reference model (red, dashed). It can be seen that fairly accurate reconstructions are obtained in spite of gender (third row) and race (top and bottom rows) differences between the faces in the input image and the reference model.

We further use the same setting to demonstrate the robustness of the algorithm. In Fig. 3 we present reconstructions of the same face using several different reference models. The face to be reconstructed differs quite significantly in shape from the reference faces due to difference in race. While there are some inaccuracies in the cheek areas, in general the recovered shapes are consistently very similar to the ground truth.

Finally, we applied the method to several real images, including some containing facial expressions. These images include one from the YaleB face database

[25], images photographed by us, and images that were downloaded from the worldwide web. For reference we used one of the 3D models from the USF database. The results are shown in Fig. 5. While we do not have the ground truth faces in these experiments, we can still see that fairly convincing reconstructions are obtained. Note in particular the reconstructions obtained with different facial expressions (right column) and the wrinkles present in the reconstruction (left column, last row).

To conclude, our experiments demonstrate that the method can accurately reconstruct faces under a large variety of uncontrolled lighting conditions and that differ from the reference face by gender, race, and expression.

## 5 Conclusion

In this paper, we have presented a novel algorithm for the recovery of 3D shape and albedo of faces from a single image by using a single reference model of different individual. Unlike existing methods, our method does not need to establish correspondence between symmetric portions of a face, nor does it require to store a database of many faces with point correspondences across the faces. Instead, our method exploits the global similarity of faces to fill in the information missing in order to apply shape recovery by solving a shape from shading problem. We tested our method by comparing the recovery obtained with rendered images to ground truth shapes and by applying the method to various real images.

Our experiments demonstrate that the method was able to accurately recover the shape of faces overcoming significant differences across individuals including differences in race, gender and variations in expressions. Furthermore we showed that the method can handle a variety of uncontrolled lighting conditions, and that it can achieve consistent reconstructions with different reference models. We hope in the future to further improve the accuracy of our method by taking an explicit account of the noise characteristics in the image and by better modeling the reflectance properties of a face (e.g., by using a second order harmonic approximation). Finally, we intend to further extend our method to handle non-frontal faces.

## References

1. Horn, B.: Obtaining Shape from Shading Information. The Psychology of Computer Vision. McGraw-Hill, New York (1975)
2. Horn, B., Brooks, M., eds.: Shape from Shading. MIT Press: Cambridge, MA (1989)
3. Rouy, E., Tourin, A.: A viscosity solutions approach to shape-from-shading. SIAM Journal of Numerical Analysis **29**(3) (1992) 867–884
4. Zhang, R., Tsai, P., Cryer, J., Shah, M.: Shape from shading: A survey. PAMI **21**(8) (1999) 690–706
5. Pentland, A.: Finding the illuminant direction. Journal Optical Society of America (1982) 448–455

6. Zheng, Q., Chellappa, R.: Estimation of illuminant direction, albedo, and shape from shading. *PAMI* **13**(7) (1991) 680–702
7. Tsai, P., Shah, M.: Shape from shading with variable albedo. *Optical Engineering* (1998) 121–1220
8. Moses, Y., Edelman, S., Ullman, S.: Generalization to novel images in upright and inverted faces. *Perception* **25** (1996) 443–461
9. Shimshoni, I., Moses, Y., Lindenbaum, M.: Shape re-construction of 3d bilaterally symmetric surfaces. *IJCV* **39**(2) (2000) 97–100
10. Zhao, W., Chellappa, R.: Symmetric shape-from-shading using self-ratio image. *IJCV* **45** (2001) 55–75
11. Atick, J., Griffin, P., Redlich, A.: Statistical approach to shape from shading: Reconstruction of three-dimensional face surfaces from single two-dimensional images. *Neural Computation* **8**(6) (1996) 1321–1340
12. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. *SIGGRAPH I* (1999) 187–194
13. Zhou, S., Chellappa, R., Jacobs, D.: Characterization of human faces under illumination variations using rank, integrability, and symmetry constraints. *ECCV* **1** (2004) 588–601
14. Smith, W., Hancock, E.: Recovering facial shape and albedo using a statistical model of surface normal direction. *ICCV* (2005) 588–595
15. Romdhani, S., Vetter, T.: Efficient, robust and accurate fitting of a 3d morphable model. *ICCV* (2003)
16. Dovgand, R., Basri, R.: Statistical symmetric shape from shading for 3d structure recovery of faces. *ECCV* (2004)
17. Sim, T., Kanade, T.: Combining models and exemplars for face recognition: An illuminating example. *CVPR Workshop on Models versus Exemplars* (2001)
18. Zhang, L., Samaras, D.: Face recognition under variable lighting using harmonic image exemplars. *CVPR I* (2003) 19–25
19. Shashua, A., Riklin-Raviv, T.: The quotient image: Class based re-rendering and recognition with varying illuminations. *PAMI* **23**(2) (2001) 129–139
20. Basri, R., Jacobs, D.: Lambertian reflectance and linear subspaces. *PAMI* **25**(2) (2003) 218–233
21. Ramamoorthi, R., Hanrahan, P.: On the relationship between radiance and irradiance: Determining the illumination from images of a convex lambertian object. *JOSA* **18**(10) (2001) 2448–2459
22. Frolova, D., Simakov, D., Basri, R.: Accuracy of spherical harmonic approximations for images of lambertian objects under far and near lighting. *ECCV* (2004) 574–587
23. Dupuis, P., Oliensis, J.: An optimal control formulation and related numerical methods for a problem in shape reconstruction. *The Annals of Applied Probability* **4**(2) (1994) 287–346
24. Kimmel, R., Sethian, J.: Optimal algorithm for shape from shading and path planning. *Journal of Mathematical Imaging and Vision* **14**(3) (2001) 237–244
25. Georghiadis, A., Belhumeur, P., Kriegman, D.: From few to many: generative models for recognition under variable pose and illumination. *PAMI* **23**(6) (2001) 643–660
26. USF DARPA Human-ID 3D Face Database: Courtesy of Prof. Sudeep Sarkar, University of South Florida, Tampa, FL. <http://marthon.csee.usf.edu/HumanID/>