Face Relighting with Radiance Environment Maps*

Zhen Wen
University of Illinois
Urbana, IL 61801
zhenwen@ifp.uiuc.edu

Zicheng Liu
Microsoft Research
Redmond, WA 98052
zliu@microsoft.com

Thomas S. Huang
University of Illinois
Urbana, IL 61801
huang@ifp.uiuc.edu

Abstract

A radiance environment map pre-integrates a constant surface reflectance with the lighting environment. It has been used to generate photo-realistic rendering at interactive speed. However, one of its limitations is that each radiance environment map can only render the object which has the same surface reflectance as what it integrates. In this paper, we present a ratio-image based technique to use a radiance environment map to render diffuse objects with different surface reflectance properties. This method has the advantage that it does not require the separation of illumination from reflectance, and it is simple to implement and runs at interactive speed.

In order to use this technique for human face relighting, we have developed a technique that uses spherical harmonics to approximate the radiance environment map for any given image of a face. Thus we are able to relight face images when the lighting environment rotates. Another benefit of the radiance environment map is that we can interactively modify lighting by changing the coefficients of the spherical harmonics basis. Finally we can modify the lighting condition of one person's face so that it matches the new lighting condition of a different person's face image assuming the two faces have similar skin albedos.

1 Introduction

Generating photo-realistic images of human faces under arbitrary lighting conditions has been a fascinating yet challenging problem. Despite its difficulty, great progress has been made in the past a few years. One class of method is the inverse rendering [16, 6, 26, 7, 8]. By capturing lighting environment and recovering surface reflectance properties, one can generate photo-realistic rendering of objects including human faces under new lighting conditions. To achieve interactive rendering speed, people have proposed to use prefiltered environment maps [18, 9]. Recently, Ramamoorthi and Hanrahan [21] used an analytic expression in terms of spherical harmonic coefficients of the lighting to approximate irradiance environment maps and they discovered that only 9 coefficients are needed for rendering diffuse objects. Basri and Jacobs [2] obtained similar theoretical results.

A method closely related to our work is the technique proposed by Cabral et al. [4]. They used pre-computed radiance environment maps to render objects rotating in the lighting environment. This technique produces photo-realistic rendering results and runs at interactive rates. One limitation of this approach is that the user has to create a radiance environment map for each different material. If the material of a surface is different at each point, it would be difficult to apply this approach. How to overcome this limitation is one of the motivations of our work. We observe that for diffuse objects, ratio image technique [23, 13] can be used to remove the material dependency of the radiance environment map. Therefore, once we capture (or compute from images as we will describe later) a radiance environment map, we can then use it to render any diffuse objects under rotated lighting environment. Note that in this paper, we only consider radiance environment map due to diffuse reflection, which is equivalent to irradiance environment map scaled by surface albedo.

In particular, we apply this technique to face images by assuming that faces are diffuse. Given a single photograph of a human face, even if we know its precise 3D geometry, it would still be difficult to separate lighting from the skin reflectance. However, it is possible to approximate the radiance environment map by using the spherical harmonics technique which was recently developed by Ramamoorthi and Hanrahan [21]. We can then use our relighting technique to generate face images under novel lighting environment. The advantage of our method is that it's simple to implement, fast and requires only one input face image.

2 Related work

In both computer vision and graphics communities, there are many works along the direction of inverse rendering. Marschner et al. [17, 15] measured the geometry and reflectance field of faces from a large number of image sam-

*Electronic version available: http://www.ifp.uiuc.edu/~ zhenwen ples in a controlled environment. Georghiades et al. [8] and Debevec et al. [7] used a linear combination of basis images to represent face reflectance. Three basis were used in [8] assuming faces are diffuse, while more basis were computed from dense image samples in [7] to handle more general reflectance functions. Sato et al. [24] and Loscos et al. [14] used ratio of illumination to modify input image to do relighting. Interactively relighting was achieved in [14] for certain point light source distributions. Ramamoorthi and Hanrahan [22] presented a signal processing framework for inverse rendering which provides analytical tools to deal with cases under general lighting conditions. They also presented practical algorithms and representation that would speed up existing inverse rendering based face relighting techniques.

Given a face under two different lighting conditions, and another face under the first lighting condition, Riklin-Raviv and Shashua [23] used color ratio (called quotient image) to generate an image of the second face under the second lighting condition. Stoschek [25] combined the quotient image with image morphing to generate relit faces under continuous changes of poses. Recently, Liu et al. [13] used the ratio image technique to map one person's facial expression details to other people's faces. One essential property of the ratio image is that it removes the material dependency thus allowing illumination terms (geometry and lighting) to be captured and transferred. In this paper, we use the ratio image technique in a novel way to remove the material dependency in the radiance environment map. As an application of this technique we generate face images under arbitrarily rotated lighting environment from a single photograph.

For the application of relighting faces from a photograph, the closest work to ours is that of Marschner et al. [16]. In their work, lighting was assumed to be a linear combination of basis lights and albedo was assumed to be known. Given the geometry, they first solve the lighting from a photograph, and then modify the lighting condition. The first difference between their work and ours is that their method is inherently an offline method because they need to integrate lighting at rendering, while our method runs at interactive rates. The second difference is that using spherical harmonics results in a more stable solver than using the linear combination of basis lights. Furthermore, it provides an effective way to interatively edit the lighting condition of a photograph. The third difference is that their theory requires known albedo, while we theoretically justify why we do not need to know the albedo when we rotate lighting. We also theoretically justify why it is feasible to assume the face has constant albedos when applying inverse rendering techniques [16].

3 Relighting with radiance environment maps

3.1 Radiance environment map

Radiance environment map was proposed by Cabral et al. [4]. Given any sphere with constant BRDF, its radiance environment map records its reflected radiance for each point on the sphere.

For any surface with the same reflectance material as that of the sphere, its reflected radiance at any point can be found from the radiance environment map by simply looking up the normal. This method is very fast and produces photorealistic results.

One limitation of this method is that one needs to generate a radiance environment map for each different material. For some objects where every point's material may be different, it would be difficult to apply this method.

To overcome this limitation, we observe that, for diffuse objects, the ratio image technique [23, 13] can be used to remove the material dependency of the radiance environment map, thus making it possible to relight surfaces which have variable reflectance properties.

Given a diffuse sphere of constant reflectance coefficient ρ, let $\mathbf{L}$ denote the distant lighting distribution. The irradiance on the sphere is then a function of normal $\vec{n}$, given by an integral over the upper hemisphere $\Omega(\vec{n})$ at $\vec{n}$.

$$E(\vec{n}) = \int_{\Omega(\vec{n})} L(\omega)(\vec{n} \cdot \omega) d\omega \tag{1}$$

The intensity of the sphere is

$$I_{sphere}(\vec{n}) = \rho E(\vec{n}) \tag{2}$$

$I_{sphere}(\vec{n})$ is the radiance environment map. Notice that radiance environment map depends on both the lighting and the reflectance coefficient ρ.

3.2 Relighting when rotating in the same lighting condition

When an object is rotated in the same lighting condition $\mathbf{L}$, the intensity of the object will change due to incident lighting changes. For any given point p on the object, suppose its normal is rotated from $\vec{n_a}$ to $\vec{n_b}$. Assuming the object is diffuse, and let ρ_p denote the reflectance coefficient at p, then the intensities at p before and after rotation are respectively:

$$I_{object}(\vec{n_a}) = \rho_p E(\vec{n_a}) \tag{3}$$

$$I_{object}(\vec{n_b}) = \rho_p E(\vec{n_b}) \tag{4}$$

From equation 3 and 4 we have

$$\frac{I_{object}(\vec{n_b})}{I_{object}(\vec{n_a})} = \frac{E(\vec{n_b})}{E(\vec{n_a})} \tag{5}$$

From the definition of the radiance environment map (equation 2), we have

$$\frac{I_{sphere}(\vec{n_b})}{I_{sphere}(\vec{n_a})} = \frac{E(\vec{n_b})}{E(\vec{n_a})} \tag{6}$$

Comparing equation 5 and 6, we have

$$\frac{I_{sphere}(\vec{n_b})}{I_{sphere}(\vec{n_a})} = \frac{I_{object}(\vec{n_b})}{I_{object}(\vec{n_a})} \tag{7}$$

What this equation says is that for any point on the object, the ratio between its intensity after rotation and its intensity before rotation is equal to the intensity ratio of two points on the radiance environment map. Therefore given the old intensity $I_{object}(\vec{n_a})$ (before rotation), we can obtain the new intensity $I_{object}(\vec{n_b})$ from the following equation:

$$I_{object}(\vec{n_b}) = \frac{I_{sphere}(\vec{n_b})}{I_{sphere}(\vec{n_a})} \cdot I_{object}(\vec{n_a}) \tag{8}$$

3.3 Comparison with inverse rendering approach

It is interesting to compare this method with the inverse rendering approach. To use inverse rendering approach, we can capture the illumination environment map, and use spherical harmonics technique [21] to obtain $E(\vec{n})$: the diffuse components of the irradiance environment map. The reflectance coefficient ρ_p at point p can be resolved from its intensity before rotation and the irradiance, that is,

$$\rho_p = \frac{I_{object}(\vec{n_a})}{E(\vec{n_a})} \tag{9}$$

After rotation, its intensity is equal to $\rho_p E(\vec{n_b})$. From equation 9 and 6,

$$\begin{aligned} \rho_p E(\vec{n_b}) &= I_{object}(\vec{n_a}) \frac{E(\vec{n_b})}{E(\vec{n_a})} \\ &= \frac{I_{sphere}(\vec{n_b})}{I_{sphere}(\vec{n_a})} \cdot I_{object}(\vec{n_a}) \end{aligned}$$

Thus, as expected, we obtain the same formula as equation 8.

The difference between our approach and the inverse rendering approach is that our approach only requires a radiance environment map, while the inverse rendering approach requires the illumination environment map (or irradiance environment map). In some cases where only limited amount of data about the lighting environment is available (such as a few photographs of some diffuse objects), it would be difficult to separate illuminations from reflectance properties to obtain illumination environment map. Our technique allows us to do image-based relighting of diffuse objects even from a single image.

3.4 Relighting in different lighting conditions

Let $\mathbf{L}$, $\mathbf{L}'$ denote the old and new lighting distributions respectively. Suppose we use the same material to capture the radiance environment maps for both lighting conditions. Let I_{sphere} and I'_{sphere} denote the radiance environment maps of the old and new lighting conditions, respectively. For any point p on the object with normal $\vec{n}$, its old and new intensity values are respectively:

$$I_{object}(\vec{n}) = \rho_p \int_{\Omega(\vec{n})} L(\omega)(\vec{n} \cdot \omega) d\omega \tag{10}$$

$$I'_{object}(\vec{n}) = \rho_p \int_{\Omega(\vec{n})} L'(\omega)(\vec{n} \cdot \omega) d\omega \tag{11}$$

Together with equation 1 and 2, we have a formula of the ratio of the two intensity values

$$\frac{I'_{object}(\vec{n})}{I_{object}(\vec{n})} = \frac{I'_{sphere}(\vec{n})}{I_{sphere}(\vec{n})} \tag{12}$$

Therefore the intensity at p under new lighting condition can be computed as

$$I'_{object}(\vec{n}) = \frac{I'_{sphere}(\vec{n})}{I_{sphere}(\vec{n})} \cdot I_{object}(\vec{n}) \tag{13}$$

4 Approximating a radiance environment map using spherical harmonics

A radiance environment map can be captured by taking photographs of a sphere painted with a constant diffuse material. It usually requires multiple views and they need to be stitched together. However, spherical harmonics technique [21] provides a way to approximate a radiance environment map from one or more images of a sphere or other types of surfaces.

Using the notation of [21], the irradiance can be represented as a linear combination of spherical harmonic basis functions:

$$E(\vec{n}) = \sum_{l \geq 0, -l \leq m \leq l} \hat{A}_l \, L_{lm} \, Y_{lm}(\vec{n}) \tag{14}$$

Ramamoorthi and Hanrahan [21] showed that for diffuse reflectance, only 9 coefficients are needed to approximate the irradiance function, that is,

$$E(\vec{n}) \approx \sum_{l \leq 2, -l \leq m \leq l} \hat{A}_l \, L_{lm} \, Y_{lm}(\vec{n}) \tag{15}$$

Given an image of a diffuse surface with constant albedo ρ, its reflected radiance at a point with normal $\vec{n}$ is equal to

$$\rho E(\vec{n}) \approx \sum_{l \leq 2, -l \leq m \leq l} \rho \, \hat{A}_l \, L_{lm} \, Y_{lm}(\vec{n}) \tag{16}$$

If we treat $\rho\, \hat{A}_l\, L_{lm}$ as a single variable for each l and m, we can solve for these 9 variables using a least square procedure, thus obtaining the full radiance environment map. This approximation gives a very compact representation of the radiance environment map, using only 9 coefficients per color channel.

An important extension is the type of surface whose albedo, though not constant, does not have low-frequency components (except the constant component). To justify this, we define a function $\rho(\vec{n})$ such that $\rho(\vec{n})$ equals to the average albedo of surface points whose normal is $\vec{n}$. We expand $\rho(\vec{n})$ using spherical harmonics as:

$$\rho(\vec{n}) = \rho_{00} + \Psi(\vec{n}) \tag{17}$$

where ρ_{00} is the constant component and $\Psi(\vec{n})$ contains other higher order components. Together with equation 16, we have

$$\begin{aligned} \rho(\vec{n})E(\vec{n}) \approx \rho_{00} & \sum_{l\leq 2, -l\leq m\leq l} \hat{A}_l\, L_{lm}\, Y_{lm}(\vec{n}) \\ + \Psi(\vec{n}) & \sum_{l\leq 2, -l\leq m\leq l} \hat{A}_l\, L_{lm}\, Y_{lm}(\vec{n}) \end{aligned} \tag{18}$$

If $\Psi(\vec{n})$ does not have first four order $(l = 1, 2, 3, 4)$ components, the second term of the righthand side in equation 18 contains components with orders equal to or higher than 3 (see Appendix for the explanation). Because of the orthogonality of spherical harmonic basis, the 9 coefficients of order $l \leq 2$ estimated from $\rho(\vec{n})E(\vec{n})$ with a linear least square procedure are $\rho_{00}\, \hat{A}_l\, L_{lm}$, where $(l \leq 2, -l \leq m \leq l)$. Therefore we obtain the radiance environment map with reflectance coefficient equal to the average albedo of the surface. This observation agrees with [22] and perception literature (such as Land's retinex theory [12]), where on Lambertian surface high-frequency variation is due to texture, and low-frequency variation probably associated with illumination.

We believe that human face skins are approximately this type of surfaces. The skin color of a person's face has dominant constant component, but there are some fine details corresponding to high frequency components in frequency domain. Therefore the first four order components must be very small. To verify this, we used *SpharmonicKit* [1] to compute the spherical harmonic coefficients for the function $\rho(\vec{n})$ of the albedo map shown in Figure 1 which was obtained by Marschner et al. [17]. There are normals that are not sampled by the albedo map, where we assigned $\rho(\vec{n})$ the mean of existing samples. We find that the coefficients of order $1, 2, 3, 4$ components are less than 6% of the constant coefficient.

5 Face relighting from a single image

Given a single photograph of a person's face, it is possible to compute its 3D geometry [3, 27]. In our implementation,

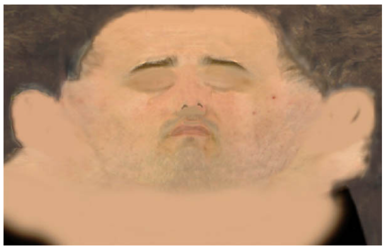

Figure 1: *A face albedo map*

we choose to use a generic geometric model (see Figure 3(a)) because human faces have similar shapes, and the artifacts due to geometric inaccuracy are not very strong since we only consider diffuse reflections.

Given a photograph and a generic 3D face geometry, we first align the face image with the generic face model. From the aligned photograph and the 3D geometry, we use the method described in section 4 to approximate the radiance environment map. To relight the face image under rotated lighting environment, we compute each face pixel's normal (with respect to the lighting environment) before and after rotation. Then we compute ratio $\frac{I_{sphere}(\vec{n_b})}{I_{sphere}(\vec{n_a})}$, where $\vec{n_b}$ and $\vec{n_a}$ are the new and old normal vectors of the pixel respectively, and I_{sphere} is the approximated radiance environment map. Finally the new intensity of the pixel is given by equation 8.

In [21] spherical harmonic basis functions of irradiance environment map were visualized on sphere intuitively, which makes it easy to modify lighting by manually changing the coefficients. Our radiance environment map is the irradiance environment map scaled by constant albedo. We can modify the coefficients in equation 16 to interactively create novel radiance environment maps and edit the lighting in the input face photograph. Unlike [21], we do not need to know the face albedo.

If we are given photographs of two people's faces under different lighting conditions, we can modify the first photograph so that it matches the lighting condition of the second face. We first compute the radiance environment map for each face. If the two faces have the same average albedo, then the two radiance environment maps have the same albedo, and we can apply equation 13 to relight the first face to match the second face's lighting condition. In practice, if two people have similar skin colors, we can apply equation 13 to relight one face to match the lighting condition of the other.

Note that it is an under-constrained problem to determine all the 9 coefficients from a single frontal image of a face according to [20]. To produce visually plausible results without conflicting with the information provided by the frontal image, we make assumptions about lighting in the back to constrain the problem. One assumption is a symmetric lighting environment, that is, the back has the same lighting distribution as the front. This assumption is equivalent to assuming $L_{lm} = 0, (l, m) = (1, 0), (2, -1), (2, 1)$ in equation 16. The rest coefficients can then be solved uniquely according to [20]. We use this assumption for lighting editing and relighting faces in different lighting conditions because only the frontal lighting matters. It also works for the rotation of the lighting environment where the lights mainly come from the two sides of the face. Otherwise, we make assumption based on the scene. For example, we assume the back is dark if the lights mostly come from the front.

5.1 Dynamic range of images

Because digitized image has limited dynamic range, ratio-based relighting would have artifacts where skin pixel values are too low or nearly saturated. To alleviate the problem, we apply constrained texture synthesis for these pixels. Our assumption is that the high frequency face albeo, similar to texture, contains repetitive patterns. Thus we can infer local face appearance at the places of artifacts from examples on other part of the face. We first identify these pixels as outliers that do not satisfy Lambertian model using robust statistics [10]. Then we use the remaining pixels as example to synthesize texture at the place of outliers. We use a patch-based Image Analogy algorithm [11], with the constraint that a candidate patch should match the original patch up to a relighting scale factor. Since we use a patch-based approach and we only apply it to the detected outlier regions, the computation overhead is very small. Figure 2 shows an example where the low dynamic range of the image causes artifacts in the relighting. (a) is the input image, whose blue channel (b) has very low intensity on the person's left face. (c) is the relighting result without using the constrained texture synthesis, and (d) is relighting result with constrained texture synthesis. We can see that almost all the unnatural color on the person's left face in (c) disappear in (d).

(a) (b) (c) (d)

Figure 2: *Using constrained texture synthesis to reduce artifacts in the low dynamic range regions. (a): input image, (b): blue channel of (a) with very low dynamic range, (b): relighting without constrained texture synthesis, and (c): relighting with constrained texture synthesis.*

6 Implementation

In our implementations, we use a cyberware-scanned face mesh as our generic face geometry (shown in Figure 3(a)). All the examples reported in this paper are produced with this mesh. Given a 2D image, to create the correspondence between the vertices of the mesh and the 2D image, we first create the correspondence between the feature points on the mesh and the 2D image. The feature points on the 2D image are marked manually as shown in Figure 3(b). We use image warping technique to generate the correspondence between the rest of the vertices on the mesh and the 2D image. A mapping from the image pixels to a radiometrically linear space is implemented to account for gamma correction. The computation of the radiance environment map takes about one second for a 640x480 input image on a PC with Pentium III 733 MHz and 512 MB memory. The face relighting is currently implemented completely in software without hardware acceleration. It runs about 10 frames per second.

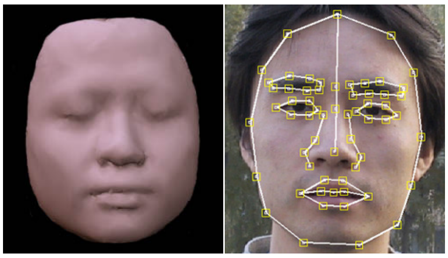

(a) (b)

Figure 3: *(a): the generic mesh. (b): feature points.*

7 Results

We show some experiment results in this section. The first example is shown in Figure 5 where the middle image is the input image. The radiance environment map is shown below the input image. It is computed by assuming the back is dark since the lights mostly come from the frontal directions. The rest of the images in the sequence show the relighting results when the lighting environment rotates. Below each image, we show the corresponding rotated radiance environment map. The environment rotates about $45°$ between each two images, a total of $180°$ rotation. The accompanying videos show the continuous face appearance changes, and the radiance environment map is shown on the righthand side of the face. (The environment rotates in such a direction that the environment in front of the person turns from the person's right to the person's left). From the middle image to the right in the image sequence, the frontal environment turns to the person's left side. On the fourth image, we can see that part of his right face gets darker. On the last image, a larger region on his right face becomes darker. This is consistent with the rotation of the lighting environment.

Figure 6 shows a different person in a similar lighting environment. For this example, we have captured the ground truth images at various rotation angles so that we are able to do a side-by-side comparison. The top row images are the ground truth images while the images at the bottom are the synthesized results with the middle image as the input. We can see that the synthesized results match very well with the ground truth images. There are some small differences mainly on the first and last images due to specular reflections (According to Marschner et al. [15], human skin is almost Lambertian at small light incidence angles and has strong non-Lambertian scattering at higher angles).

The third example is shown in Figure 7. Again, the middle image is the input. In this example, the lights mainly come from two sides of the face. The bright white light on the person's right face comes from sky light and the reddish light on his left face comes from the sun light reflected by a red-brick building. The image sequence shows a 180° rotation of the lighting environment.

Figure 4 shows four examples of interactive lighting editing by modifying the spherical harmonics coefficients. For each example, the left image is the input image and the right image is the result after modifying the lighting. In example (a), lighting is changed to attach shadow on the person's left face. In example (b), the light on the person's right face is changed to be more reddish while the light on her left face becomes slightly more blueish. In (c), the bright sunlight move from the person's left face to his right face. In (d), we attach shadow to the person's right face and change the light color as well. Such editing would be difficult to do with the currently existing tools such as *Photoshop*.

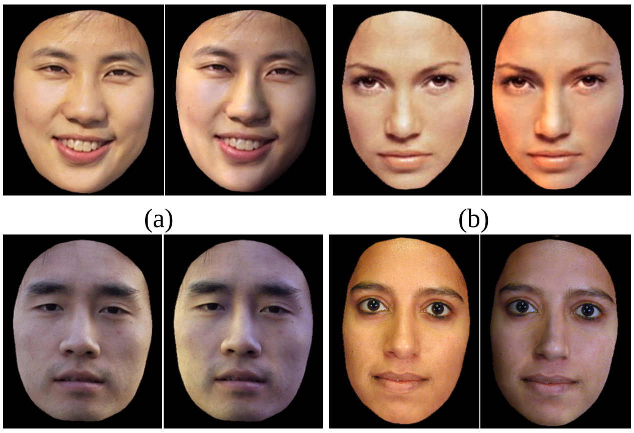

(a) (b) (c) (d)

Figure 4: *Interactive lighting editing by modifying the spherical harmonics coefficients of the radiance environment map.*

We have also experimented with our technique to relight a person's face to match the lighting condition on a different person's face image. As we pointed out in Section 5, our method can only be applied to two people with similar skin colors. In Figure 8(a), we relight a female's face shown on the left to match the lighting condition of a male shown in the middle. The synthesized result is shown on the right. Notice the darker region on the right face of the middle image. The synthesized result shows similar lighting effects. Figure 8(b) shows an example of relighting a male's face to

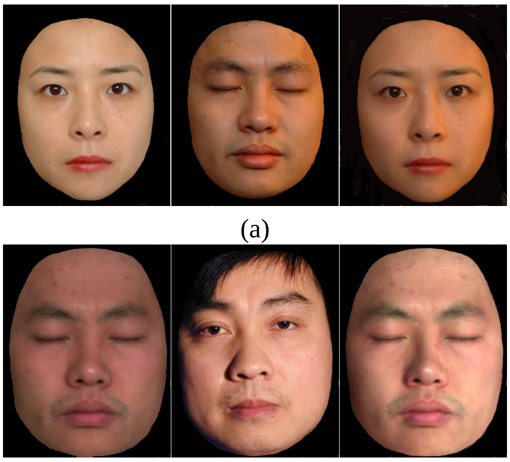

(a) (b)

Figure 8: *Relighting under different lighting. For both (a) and (b): Left: Face to be relighted. Middle: target face. Right: result.*

a different male. Again, the left and middle faces are input images. The image on the right is the synthesized result. From the middle image, we can see that the lighting on his left face is a lot stronger than the lighting on his right face. We see similar lighting effects on the synthesized result. In addition, the dark region due to attached shadow on the right face of the synthesized image closely matches the shadow region on the right face of the middle image.

8 Conclusion and future work

We have presented a novel technique that, given a single image of a face, generates relit face images when the lighting environement rotates. We have also showed that the computed radiance environment map provides a new way of editing lighting effects by adjusting the sperical harmonics coefficients. Finally we can relight one person's face image to match the new lighting condition of a different person's face image if the two people have similar skin colors. Experiments show that our method is effective and works in real time.

We are planning on developing a more intuitive user interface for interactively editing lighting effects. We hope that our system will allow the user to use higher level commands such as making this area brighter, more red on that area, etc. Currently the Lambertian model used in our work can not deal with specular reflection, subsurface scattering on human faces. In the future, we would like to extend our methods to handle more general non-Lambertian BRDFs based on the theoretical results in [22]. One possible way

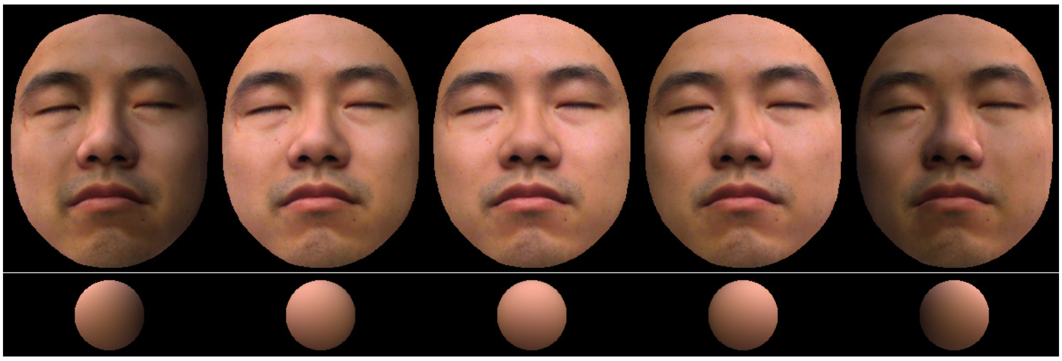

Figure 5: *The middle image is the input. The sequence shows synthesized results of* 180° *rotation of the lighting environment.*

Figure 6: *The comparison of synthesized results and ground truth. The top row is the ground truth. The bottom row is synthesized result, where the middle image is the input.*

to deal with specular reflection is to separate specular from diffuse reflections such as the method in [19]. Our method also neglects cast shadows and inter-reflections on face images, as is common in radiance map algorithms. To handle those effects, techniques using more basis images such as that in [7] will be useful. The difference between the 3D generic face geometry and the actual one may introduce artifacts. We are planning on estimating a personalized geometric model from the input image by using techniques such as those reported in [3, 27]. We also would like to apply our technique to other types of objects.

Acknowledgments

This work was supported in part by National Science Foundation Grant CDA 96-24386, and IIS-00-85980. Thanks Brian Guenter for the face albedo map.

Appendix

The multiplication of two spherical harmonic basis satisfies the following relation [5]:

$$Y_{l_1 m_1}(\vec{n}) Y_{l_2 m_2}(\vec{n}) = \sum_{l=|l_1-l_2|}^{l_1+l_2} \sum_{m=-l}^{l} \{ C\langle l_1, l_2 : 0, 0 | l, 0 \rangle \cdot C\langle l_1, l_2 : m_1, m_2 | l, m \rangle Y_{lm}(\vec{n}) \} \quad (19)$$

where $C\langle\, l_1, l_2 : m_1, m_2 \,|\, l, m \,\rangle$ is the Clebsch-Gordan coefficients. The coefficient of $Y_{lm}(\vec{n})$ in the righthand side is non-zero if and only if $m = m_1 + m_2$, l range from $|l_1 - l_2|$ to $l_1 + l_2$ and $l_1 + l_2 - l$ is even.

We then look at the second term of the righthand side in equation 18. Suppose $Y_{l_1\, m_1}(\vec{n})$ comes from $\Psi(\vec{n})$, and $Y_{l_2\, m_2}(\vec{n})$ comes from the factor corresponding to lighting. For any $Y_{l_2\, m_2}(\vec{n})$, we have $0 \le l_2 \le 2$. If $l_1 > 4$ for any $Y_{l_1\, m_1}(\vec{n})$, we have $l \ge |l_1 - l_2| > 2$. That proves the claim that if $\Psi(\vec{n})$ does not have the first four order ($l = 1, 2, 3, 4$) components, the second term of the righthand side

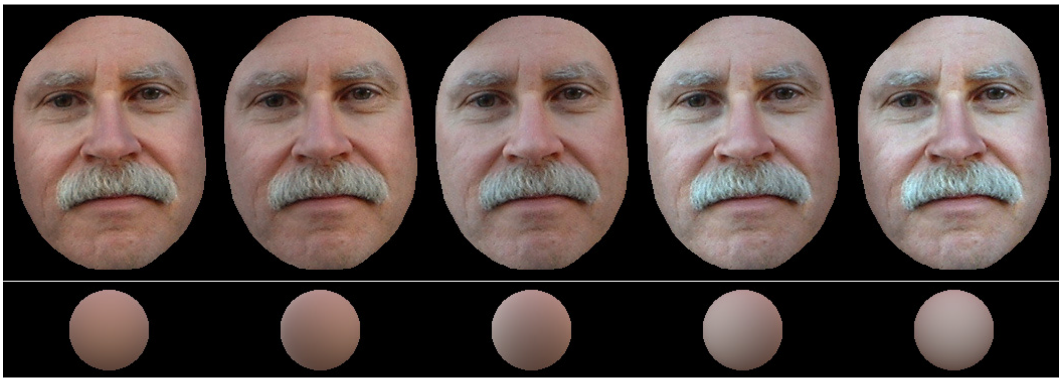

Figure 7: *The middle image is the input. The sequence shows a* 180° *rotation of the lighting environment.*

in equation 18 contains components with orders equal to or higher than 3.

References

[1] http://www.cs.dartmouth.edu/∼geelong/sphere/.

[2] R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. In *Proc. ICCV'01*, pages 383–390, 2001.

[3] V. Blanz and T. Vetter. A morphable model for the synthesis of 3d faces. In *Proc. SIGGRAPH 99*, pages 187–194, 1999.

[4] B. Cabral, M. Olano, and P. Nemec. Reflection space image based rendering. In *Proc. SIGGRAPH 99*, pages 165–170, August 1999.

[5] C. Cohen-Tannoudji and et al. *Quantum Mechanics*. John Wiley & Sons, 1977.

[6] P. E. Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proc. SIGGRAPH 98*, pages 189–198, July 1998.

[7] P. E. Debevec and *et. al*. Acquiring the reflectance field of a human face. In *Proc. SIGGRAPH 2000*, pages 145–156, July 2000.

[8] A. Georghiades, P. Belhumeur, and D. Kriegman. Illumination-based image synthesis: Creating novel images of human faces under differing pose and lighting. In *IEEE Workshop on Multi-View Modeling and Analysis of Visual Scenes*, pages 47–54, 1999.

[9] N. Greene. Environment mapping and other applications of world projections. *IEEE Computer Graphics and Applications*, 6(11):21–29, 1986.

[10] F. R. Hampel and *et. al*. *Robust Statistics*. John Wiley & Sons, 1986.

[11] A. Hertzmann and *et. al*. Image analogies. In *Proc. SIGGRAPH 2001*, pages 327–340, August 2001.

[12] E. Land and J. McCann. Lightness and retinex theory. *Journal of the Optical Society of America*, 61(1):1–11, 1971.

[13] Z. Liu, Y. Shan, and Z. Zhang. Expressive expression mapping with ratio images. In *Proc. SIGGRAPH 2001*, pages 271–276, 2001.

[14] C. Loscos, G. Drettakis, and L. Robert. Interactive virtual relighting of real scenes. *IEEE Trans. on Visualization and Computer Graphics*, 6(3), February 2000.

[15] S. R. Marschner and et al. Image-based brdf measurement including human skin. In *Rendering Techniques*, 1999.

[16] S. R. Marschner and D. P. Greenberg. Inverse lighting for photography. In *IST/SID Fifth Colort Imaging Conference*, November 1997.

[17] S. R. Marschner, B. Guenter, and S. Raghupathy. Modeling and rendering for realistic facial animation. In *Rendering Techniques*, pages 231–242. Springer Wien New York, 2000.

[18] G. Miller and C. Hoffman. Illumination and reflection maps: Simulated objects in simulated and real environments. In *SIGGRAPH 84 Advanced Computer Graphics Animation Seminar Notes*. Siggraph, 1984.

[19] K. Nishino, Z. Zhang, and K. Ikeuchi. Determining reflectance parameters and illumination distribution from a sparse set of images for view-dependent image synthesis. In *Proc. ICCV'01*, pages 599–606, 2001.

[20] R. Ramamoorthi. Analytic pca construction for theoretical analysis of lighting variability in images of a lambertian object. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 24(10), 2002.

[21] R. Ramamoorthi and P. Hanrahan. An efficient representation for irradiance environment maps. In *Proc. SIGGRAPH 2001*, pages 497–500, August 2001.

[22] R. Ramamoorthi and P. Hanrahan. A signal-processing framework for inverse rendering. In *Proc. SIGGRAPH 2001*, pages 117–128, August 2001.

[23] T. Riklin-Raviv and A. Shashua. The quotient image: Class based re-rendering and recongnition with varying illuminations. In *Proc. CVPR*, pages 566–571, June 1999.

[24] I. Sato, Y. Sato, and K. Ikeuchi. Acquiring a radiance distribution to superimpose virtual objects onto a real scene. *IEEE Trans. on Visualization and Computer Graphics*, 5(1):1–12, February 1999.

[25] A. Stoschek. Image-based re-rendering of faces for continuous pose and illumination directions. In *Proc. CVPR*, pages 582–587, 2000.

[26] Y. Yu and et al. Inverse global illumination: Recovering reflectance models of real scenes from photogrpahs. In *Proc. SIGGRAPH 99*, pages 215–224, July 1999.

[27] L. Zhang and et al. Single view modeling of free-form scenes. In *Proc. CVPR*, pages I990–I997, 2001.