



Image-based Face Illumination Transferring using Logarithmic Total Variation Models*

Qing Li, Wotao Yin[†], and Zhigang Deng

Department of Computer Science
University of Houston
Houston, TX, 77204, USA
<http://www.cs.uh.edu>

Technical Report Number UH-CS-09-06

May 7, 2009

Keywords: Illumination Transferring, Face Relighting, Logarithmic Total Variation Model, Radial Basis Functions, Face Decomposition

Abstract

In this paper, we propose a novel image-based technique that transfers illumination from a source face image to a target face image based on the Logarithmic Total Variation (LTV) model. Our method does not require any prior information regarding the lighting conditions or the 3D geometries of the underlying faces. We first use a Radial Basis Functions (RBFs)-based deformation technique to align key facial features of the reference 2D face with those of the target face. Then, we employ the LTV model to factorize each of the two aligned face images to an illumination-dependent component and an illumination-invariant component. Finally, illumination transferring is achieved by replacing the illumination-dependent component of the target face by that of the reference face. We tested our technique on numerous grayscale and color face images from various face datasets including the Yale face Database, as well as the application of illumination-preserved face coloring.



*It has been accepted to The Visual Computer in May 2009.

[†]Wotao Yin is with Department of Computational and Applied Mathematics, Rice University, Houston, TX.

Image-based Face Illumination Transferring using Logarithmic Total Variation Models*

Qing Li, Wotao Yin[†], and Zhigang Deng

Abstract

In this paper, we propose a novel image-based technique that transfers illumination from a source face image to a target face image based on the Logarithmic Total Variation (LTV) model. Our method does not require any prior information regarding the lighting conditions or the 3D geometries of the underlying faces. We first use a Radial Basis Functions (RBFs)-based deformation technique to align key facial features of the reference 2D face with those of the target face. Then, we employ the LTV model to factorize each of the two aligned face images to an illumination-dependent component and an illumination-invariant component. Finally, illumination transferring is achieved by replacing the illumination-dependent component of the target face by that of the reference face. We tested our technique on numerous grayscale and color face images from various face datasets including the Yale face Database, as well as the application of illumination-preserved face coloring.

Index Terms

Illumination Transferring, Face Relighting, Logarithmic Total Variation Model, Radial Basis Functions, Face Decomposition

I. INTRODUCTION

Image-based illumination transferring (*i.e.*, face relighting) has attracted growing attention in computer graphics and computer vision community in recent years. The face illumination transferring problem is particularly challenging when little information about the source or the target face is available, for example, only two face images (one for the source/reference face and the other for the target face) are provided, and prior lighting/illumination conditions of the faces and their 3D face geometries are not available.

In this paper, we propose a novel image-based face illumination transferring technique that transfers illumination between two different 2D faces without requiring prior information of the faces, *e.g.*, lighting conditions or 3D face geometries. As shown in Figure 1, we first warp the source (reference) face image using a Radial Basis Functions (RBFs)-based deformation technique to align its key facial features with those of the target face image. Then, we apply the Logarithmic Total Variation (LTV) model [1], [2] to factorize each of the two aligned face images into an illumination-dependent component u and an illumination-invariant component v , respectively. Finally, we relight the target face by replacing its illumination-dependent component u with that of the aligned reference face.

Compared with previous work [3]–[6], our face illumination transferring technique has the following distinctions:

- *No prior information requirement.* Our technique does not require any prior information regarding the lighting conditions or the 3D geometries of the underlying faces.
- *Robustness and efficiency.* By testing our technique on various face datasets, we found that our technique is robust, and the relighted results are close to ground truths.
- *Versatility.* Besides transferring illumination from one grayscale (or color) face to another, our approach can colorize a grayscale 2D face image, based on a reference color face image. In addition, our approach can also be used to transfer illumination between 2D side-view face images.

The remainder of this paper is organized as follows: Section II reviews related work in image-based face relighting and illumination transferring, Section III describes how to use Radial Basis Functions (RBFs) based deformation techniques to align two face images, Section IV describes the reflectance model used in this work, Section V details the Logarithmic Total Variation (LTV) model, and Section VI describes the illumination transferring algorithm. Finally, experimental results (Section VII) and conclusions (Section VIII) are presented.

*It has been accepted to The Visual Computer in May 2009.

[†]Wotao Yin is with Department of Computational and Applied Mathematics, Rice University, Houston, TX.

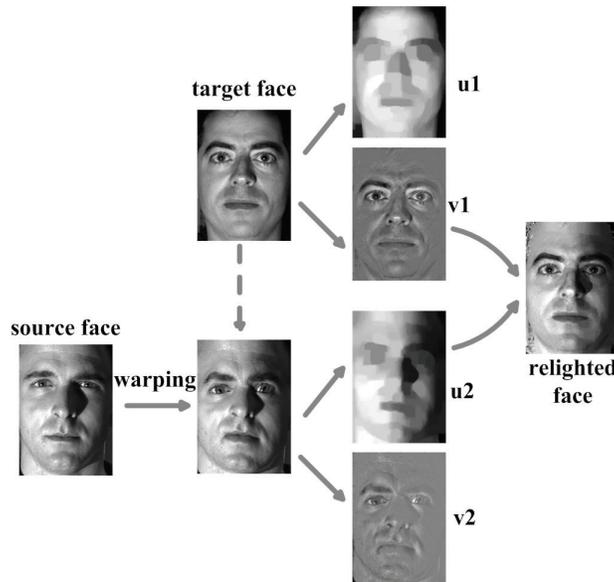


Fig. 1. The schematic overview of our algorithm. First, we employ a geometric deformation technique to align key facial features of input faces. Then each face is factorized into an illumination-dependent component u and an illumination-free component v . Finally, by switching the above components, we generate the relighted face.

II. RELATED WORK

Shashua and Raviv [7] introduce the idea of Quotient (ratio) image for face relighting. In their approach, a face is considered as a Lambertian object, and three face images are needed to generate the quotient image through a linear combination operation. The requirement of a few example images of the target face limits the use of this approach in many cases, and their approach can only work on the faces under a fixed viewpoint. Stoschek [3] further extends the Shashua-Raviv’s work to more general cases including relighting faces under several different viewpoints or wearing glasses. Liu *et al.* [8] use the quotient image to transfer fine facial details (*e.g.*, facial creases and wrinkles) to expressive faces. Peers *et al.* [5] extend the face relighting from a still image to an image sequence, where an optical flow algorithm is used to track and warp the quotient images of in-between frames. However, the basic idea of these approaches is the same: *using the quotient image for relighting a still face image*.

Other approaches were also proposed for face relighting. Marschner and Greenberg [9] introduce inverse lighting to modify the lighting of existing face images. However, the success of their approach mostly relies on the 3D surface model of the face that is difficult to acquire without considerable efforts in practice. Wen *et al.* [4] relight faces by building a radiance environment map approximation. This approach uses a spherical harmonics based environment map to capture facial variations due to diffuse lightings, and uses a ratio image to remove material dependency. Recently, Wang *et al.* [6] proposed a novel method to change the illumination condition of a single face image without knowing its face geometry and albedo information, by using a Markov Random Field to model the statistical distribution and the spatial coherence of face texture, especially for harsh lighting conditions. In addition, Chen *et al.* [10] demonstrated the success of applying the Logarithmic Total Variation (LTV) model [2] for face recognition under varying illuminations. Their approach removes illuminations from 2D face images using the LTV model, which results in a higher face recognition accuracy.

III. FACE ALIGNMENT

Before illumination can be transferred between a source face image I^s and a target face image I^t , we need to align the source (reference) face with the target face image by deforming I^s such that its key facial features including eyes, chins, nose, etc., in the deformed source face I^{ds} are aligned with those of I^t . Radial Basis Functions (RBFs) has been successfully applied to 3D geometric deformation [11]–[14] and image warping [15]. Here we briefly describe how to use Radial Basis Functions to deform the source (reference) face image.

A radial basis function $\bar{h}(x)$ is a real-valued function whose value depends only on $\|x\|$ or $\|x-c\|$ for certain center point c ; in other words, a radial basis function \bar{h} has the representation $\bar{h}(x) = h(\|x\|)$ or $\bar{h}(x, c) = h(\|x-c\|)$, where

h is some scalar function. Among the most common radial basis functions are the multi-quadrics $h(r) = \sqrt{r^2 + c^2}$, the Gaussian functions $h(r) = e^{-\frac{r}{c}}$, and the thin plate spline $h(r) = r^2 \log(r)$. In our application, we let $x \in \mathbb{R}^2$ and $c_i \in \mathbb{R}^2, \forall i$, denote an arbitrary point and a set of given feature points, respectively. We use the thin plate splines to approximate the face wrapping mapping $F: \mathbb{R}^2 \rightarrow \mathbb{R}^2$:

$$F(x) = \sum_{i=1}^n w_i h(\|x - c_i\|), \quad (1)$$

where the weights $w_1, \dots, w_n \in \mathbb{R}^2$ are determined as described below.

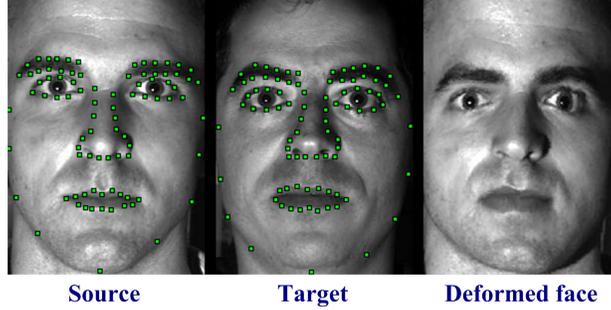


Fig. 2. The left face is the source face with desired light condition, the middle face is the target face, the right is the deformation result of the source face.

First, we select a set of $2n$ feature points in the source and target faces in pairs. Each pair consists two points at the same facial feature location, one in each of the two faces. Figure 3 depicts a set of feature points, most of which lie on the boundaries of facial objects, as well as the entire face. Let the n points in the source face be denoted $\{x_i^s\}$ and the n points in the target face by $\{x_i^t\}$. In (1), we let $c_i \equiv x_i^t$ for all i and use $F(x)$ wrap the source face to the target face. Therefore, F should satisfy

$$F(x_j^s) = x_j^t, \quad \forall j = 1, \dots, n. \quad (2)$$

Plugging (1) and the definition of thin plate spline into all of Eq. (2) yields the following linear equation system:

$$x_j^t = \sum_{i=1}^n w_i \|x_i^s - x_j^s\|^2 \log(\|x_i^s - x_j^s\|), \quad \forall j. \quad (3)$$

We solve (3) for the unknowns w_i to determine the wrapping mapping F . Next, the entire source face I^s is deformed into I^{ds} by applying F ; that is, $I^{ds} \leftarrow F(I^s)$, point-wise. Figure 2 shows an example of RBFs-based 2D deformation for face alignment.

Automatically identifying key facial features from images is a non-trivial task. A number of automated algorithms have been proposed [16], [17]. For example, Zhou *et al.* [17] proposed a Bayesian Tangent Shape Model (BTSM) for facial feature identification. Since this work focuses on face illumination transferring, not automated face alignment, in this step we choose to manually label key facial feature points on the face images due to its higher accuracy.

IV. REFLECTANCE MODEL

The Logarithm Total Variation (LTV) model is based on a general multiplicative light reflectance model:

$$I(x) = \rho(x)S(x) \quad (4)$$

where $I(x)$ is the intensity of *reflected* light at x , $\rho(x)$ is the surface albedo at x , and $S(x)$ is the amount of light received at x that generates the observed intensity. It generalizes the well-known Lambertian model, $I = A\rho \cos \theta$, where A is the strength of the point light source and θ is the angle between the light source direction and the surface normal, by letting $S = A \cos \theta$. The LTV model does not obtain ρ and S from an input I but factorizes I into $v = \rho_v$ and $u = \rho_u S$, i.e.,

$$I(x) = v(x)u(x) = \rho_v(x)(\rho_u(x)S(x)), \quad (5)$$

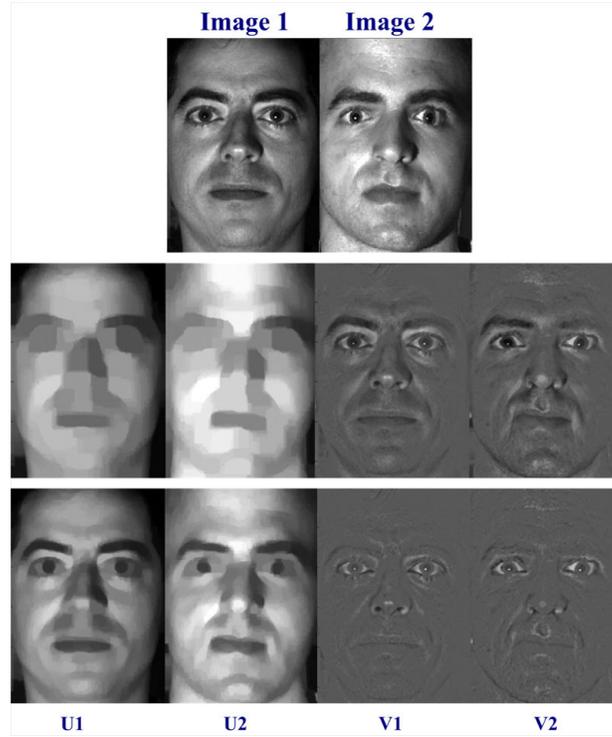


Fig. 3. Illustration of how the λ parameter controls level of details preserved in the v component. The first row is two input images. The other two rows (from top to bottom) are the output results when $\lambda = 0.2, 0.5$, respectively. The left two column are the u components of the two input images, and the right two columns are the v components of the input images.

so v represents the intensity *variations* of facial features, which are precisely the contours and boundaries of facial features such that eyes, eyebrows, lips, chins, noses, nostrils, etc. These intensity variations are often sufficient to characterize a face, and they are lighting-invariant. Because two faces may also differ in complexion, relighting a target face requires us to assume that the target face has complexion similar to that of the source face. Under this assumption, a target face can be relighted by replacing its u -component by that of the source face.

Next, we describe the TV- L^1 model, from which we obtain the multiplicative decomposition (5).

V. THE TV- L^1 MODEL

The TV- L^1 model [2] is an *additive* image model, i.e., given an input image f , it generates the decomposition $f = \tilde{u} + \tilde{v}$. To obtain the multiplicative LTV decomposition, we substitute f , \tilde{u} , and \tilde{v} in the TV- L^1 model by $\log(I)$, $\log(u)$, and $\log(v)$, respectively, and get $\log(I) = \log(u) + \log(v)$, which is equivalent to $I = u \circ v$. Next, we argue that the TV- L^1 model separates f into a large-scale, illumination-dependent component \tilde{u} and a small-scale, illumination-invariant component \tilde{v} .

The TV- L^1 model solves the optimization problem

$$\min_{\tilde{u}} \int |\nabla \tilde{u}| + \lambda \|f - \tilde{u}\|_{L^1}, \quad (6)$$

where f is the input image and \tilde{u} is the unknown and output image. The first term $\int |\nabla \tilde{u}|$ in (6) is the total variation (TV) of \tilde{u} , and λ is a scalar, which has the function as a threshold on scale. Solving (6) is equivalent to finding the solution to the following PDE:

$$\nabla \cdot \left(\frac{\nabla \tilde{u}}{|\nabla \tilde{u}|} \right) + \lambda \frac{f - \tilde{u}}{|f - \tilde{u}|} = 0,$$

where we assume $0/0 = 0$. Similar to other TV-based models, the TV- L^1 model keeps the sharp object edges in u because minimizing the TV regularization term penalizes all intensity variations in a uniform way (in contrast to $\min \int |\nabla \tilde{u}|^2$, which penalizes larger changes, like sharp edges, more heavily.) This property is very important in illumination normalization since the sharp boundaries of shadows cast on faces are entirely preserved in \tilde{u} and,

therefore, such boundaries do not affect the recognition process. On the other hand, any image smoothing models that do not preserve edges well enough will inevitably cause *halo* artifacts and affect face relighting results.

What also distinguishes the TV- L^1 model from other TV-based models is its unique properties: (i) scale dependent, but intensity-independent, decomposition, and (ii) easy parameter selection. These properties are very important to face decomposition. Let us first explain this property by example.

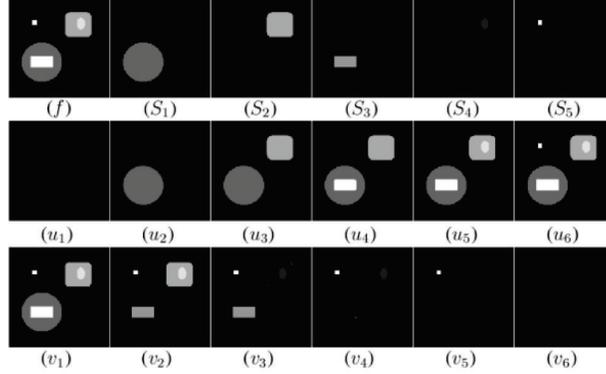


Fig. 4. The TV- L^1 decompositions of f into $u_i + v_i$ obtained using $(\lambda_1, \dots, \lambda_6) = (0.0515, 0.0746, 0.1256, 0.2188, 0.4263, 0.6000)$.

The results depicted in Figure 4 (u_1 - u_6 , v_1 - v_6) were obtained by applying the TV- L^1 model with different λ 's to the composite input image depicted in Figure 4 (f). The five components in this composite image are depicted in Figure 4 ($S_1 - S_5$) in the descending order of their scales, which are the image features that we shall extract from f . We chose λ_i , $i = 1, \dots, 6$, according to

$$\lambda_1 < \bar{\lambda}_1 < \lambda_2 < \bar{\lambda}_2 < \dots < \bar{\lambda}_5 < \lambda_6, \quad (7)$$

where each $\bar{\lambda}_i$ was computed solely based upon the shape of S_i (hence, the intensity or location of S_i does not affect $\bar{\lambda}_i$.) Specifically, we let $\bar{\lambda}_i = 1/G(\partial|\mathbf{1}_{S_i}|)$, where $G(\cdot)$ stands for the so-called G -value (refer to [2] for its mathematical and computational detail). Then, we applied the TV- L^1 model to f using $\lambda = \lambda_1, \dots, \lambda_6$, whose values are given in the caption of Figure 4. Instead of these values, we could instead use any λ_i 's satisfying (7) and still obtain the same results. The numerical results depicted in Figure 4 closely match the analytic results – λ can accurately control the selection of geometric features with arbitrary intensity values and at arbitrary locations. Hence, selecting an appropriate λ for the small-scale facial features is straightforward. We shall pick a λ that is slightly smaller than $1/G(\partial|\mathbf{1}_M|)$, where M is the facial feature mask set.

In our experiment, we found that setting λ to 0.2 worked well. Figure 3 shows how λ controls the level of details retained in the v component in the multiplicative model (5). Finally, we note that in practice, one does not need to compute G -value. A λ -value obtained by trial and error will be good for all faces of a similar size, and this value should vary inversely proportional to the face size.

VI. FACE ILLUMINATION TRANSFERRING ALGORITHM

Let I^{ds} and I^t denote the deformed source and the target faces, respectively. The factorizations $I^{ds} = u^{ds} \circ v^{ds}$ and $I^t = u^t \circ v^t$ described in Section V can be obtained by solving Problem (6) for $(f, \tilde{u}) = (\log(I^{ds}), \log(u^{ds}))$ and $(\log(I^t), \log(u^t))$, respectively. To transfer the lighting condition of I^{ds} to I^t , we use u^{ds} to replace u^t and obtain the relighted face as $I^r = u^{ds} \circ v^t$. This algorithm is described in Algorithm 1.

It should be noted that if input is a color face, Algorithm 1 needs to be executed three times, one time for each color channel (red, green, or blue).

There are various numerical algorithms for solving (6). In this work, we used the algorithm proposed by Goldfarb and Yin [18] that is based on decomposing (6) into multiple subproblems and solving all of these subproblems by a single call to a parametric max-flow algorithm. We found that this approach performed well in terms of both computational efficiency and solution accuracy. Applied to a pair of 10×5 face images, the entire process took only 3.24 seconds on a PC with 3.20GHz CPU and 1GB of memory.

Algorithm 1 The Face Illumination Transferring Algorithm

Input: a source face I^s , a target face I^t
Output: the relighted target face I^{rt}

- 1: $I^{ds} \leftarrow$ deform I^s to match I^t (using RBF)
 - 2: $u^{ds} \leftarrow \exp(\tilde{u})$, where \tilde{u} solves (6) for $f = \log(I^{ds})$
 - 3: $v^t \leftarrow \exp(f - \tilde{u})$, where \tilde{u} solves (6) for $f = \log(I^t)$
 - 4: return the relighted face $I^{rt} \leftarrow u^{ds} \circ v^t$
-

VII. EXPERIMENT RESULTS AND DISCUSSION

We conducted various experiments to test and validate our algorithm. Most of the face images used in our experiments were acquired from the Yale Face Database B [19], the CVL Face Database [20], and the AR Face Database [21].

First, ground-truth comparison experiments were conducted to validate our algorithm. We chose face images of four different subjects under various lighting conditions from the Yale Face Database B [19]. The Yale Face Database is the only database that we found containing the ground truths of every subject under various lighting conditions. We transferred the face illumination between every two subjects and compared our synthesized (relighted) results with the ground truths. Statistically, the average MAE (Mean Absolute Error) between our results and the ground truths is 8.69%. Visually, compared with the ground truths, the shapes and locations of the shadow on the faces (the relighted and the ground-truths) are very close even though nuance differences still exist. Figure 5 and Figure 6 show two examples of these comparison experiments.

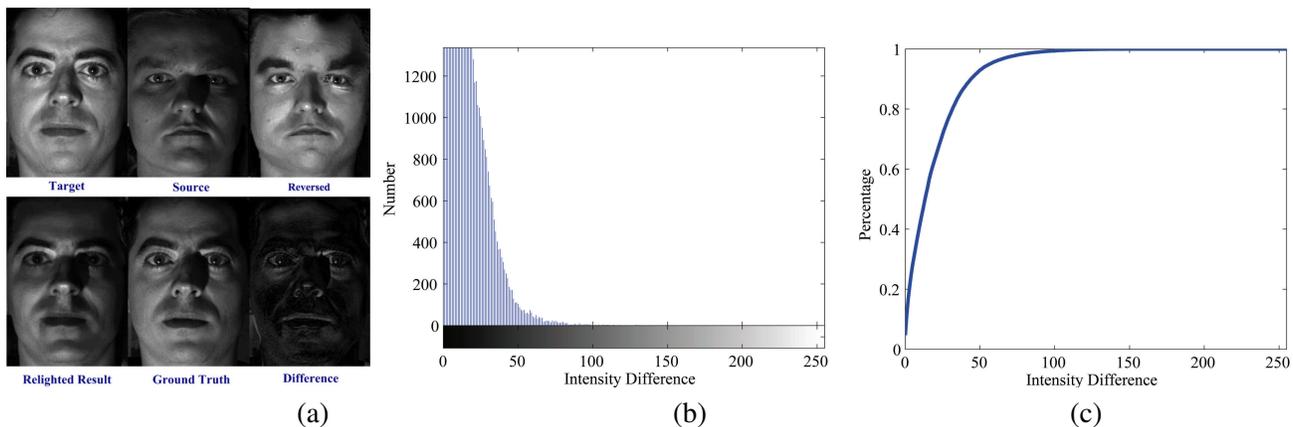


Fig. 5. The first ground truth comparison and validation example: (a) The top left is the target face, the top middle shows the source (reference) face with the desired lighting condition, the top right shows the relighted result when we exchange the roles of the target and the source face, the bottom left is the relighted face, the bottom middle shows the ground truth, and the bottom right shows the difference between the relighted face and the ground truth. The MAE in this case is 6.63%. Both the source and the target images were taken from the Yale Face Database B. (b) The histogram of the pixel-intensity difference between the relighted face and the ground truth. (c) The accumulated percentage of (b).

Besides the above grayscale face images, we also tested illumination transferring between color face images. Due to the lack of ground truths, we can not perform ground-truth comparison and statistical analysis. Figures 7 and 8 show some relighted results between color face images from both front and side views. As shown in Figures 7 and 8, although the appearances (face shape, skin reflectance, etc.) and lightings of the two human faces are very different, the relighted faces are visually plausible and perceptually realistic.

We also applied our approach in a novel application (termed *illumination-preserved face coloring*) that automatically transfers illumination from existing color face images to grayscale ones. In this experiment, we replace the u -components of the target (grayscale) faces with those of the reference (color) faces. Figure 9 depicts two such examples.

It is noteworthy that although our algorithm works well for most of test face images, certain visual artifacts can not be completely avoided in certain illumination transferring cases. However, our approach provides a user-control

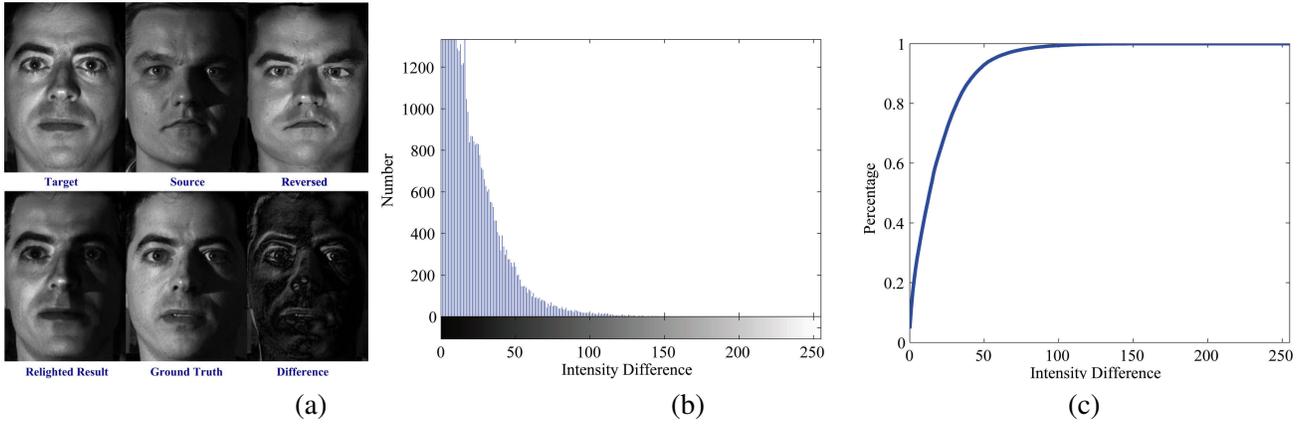


Fig. 6. The second ground truth comparison and validation example: (a) The top left is the target face, the top middle shows the source (reference) face with the desired lighting condition, the top right shows the relighted result when we exchange the roles of the target and the source face, the bottom left is the relighted face, the bottom middle shows the ground truth, and the bottom right shows the difference between the relighted face and the ground truth. The MAE in this case is 7.54%. Both the source and the target images were taken from Yale Face Database B. (b) The histogram of the pixel-intensity difference between the relighted face and the ground truth. (c) The accumulated percentage of (b).

mechanism through adjusting the λ parameter to reduce or even diminish some of these artifacts.

Using our approach, a face image can be decomposed into an illumination-dependent component and an illumination-invariant component. Ideally, the illumination-dependent component should not contain any individual-dependent information. However, certain individual-dependent information, such as individualized complexion information and certain shading effects caused by the wrinkles or any tiny movement of the subject, is hard to be completely separated from the illumination-dependent component. As we mentioned in Section V, the parameter λ controls the level of details preserved in the illumination-invariant component v . We need to exploit and tune this control parameter in order to separate the individual dependent information out of the illumination-dependent component. Figure 10 shows an example where adjusting the λ parameter substantially affects the relighting results. In this example, the complexions of the target face and the source face are very different, and certain shading effect caused by the wrinkle appears around the corners of the mouth in the source face. As we can see in Figure 10, by decreasing λ , more individualized details are retained in the illumination-invariant component v of the target face, and the illumination-dependent information (preserved in the u component of the source face) becomes more and more generalized. Therefore, when λ is decreased, the complexion in the relighted result is more closer to that of the target face, and the shading artifact is gradually reduced.

Another potential cause of relighting artifacts is the uneven illumination on the target face. Typically, a dark shade can be observed on the face under such illumination condition. During the face decomposition, the edge of the shadow in the target face may be preserved in the v component as the illumination-invariant information; and eventually, this will cause artifacts on the relighted face. Tuning the parameter λ would be able to reduce this type of visual artifacts as well. Fig 11 illustrates such an example. The more λ is increased, the less the edge of the shade is preserved in its v component. Therefore, less artifacts can be observed on the relighted face. However, we need to be cautious that a larger λ may lead to the visual loss of some other face details such as those around the eyebrows.

We also tested our algorithm in color spaces other than the RGB space. Basically, we first converted pixel colors from the RGB space to the new color space such as the HSV and YCbCr spaces, and then performed illumination transferring in the new space, and finally, an inverse color space transformation was performed. We found that if illuminations were transferred in all three color channels in the HSV or YCbCr case, the resultant faces were almost identical to the RGB-space results. If illumination was transferred only in the primary channel of the new color space (e.g., the H channel in the HSV space or the Y channel in the YCbCr space), the speed of our algorithm can be improved three times. However, the relighting results are not visually close to the expected. Figure 12 shows such an example.

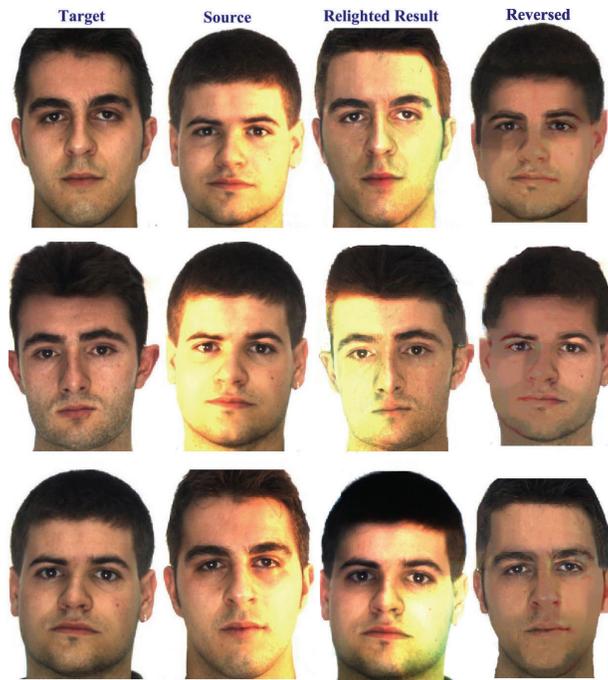


Fig. 7. Results of illumination transferring between two color faces: the first column contains the target faces, the second column contains the source (reference) faces with the desired lighting conditions, the third column contains the relighted faces, and the fourth columns contains the relighted results when the roles of the target and the source face are switched.

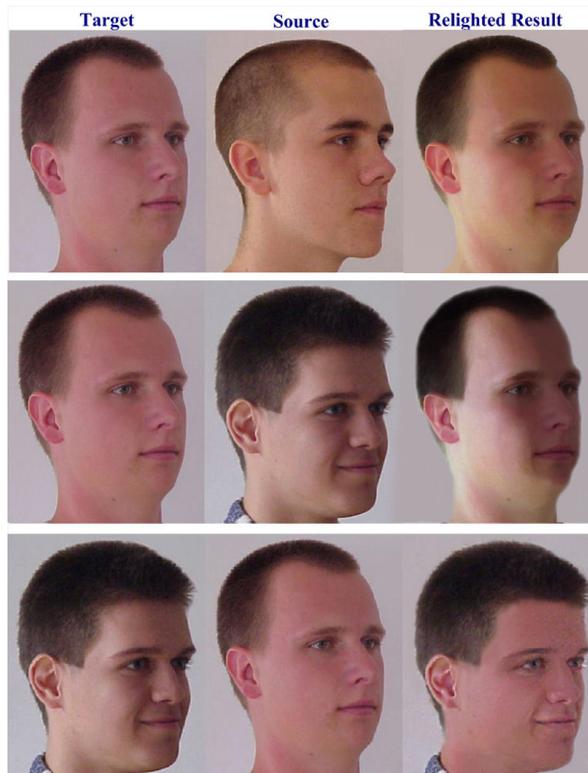


Fig. 8. Results of illumination transferring between two side-view color faces. Left column: the target faces; middle column: the source (reference) faces with desired lighting conditions; and right column: the relighted faces.

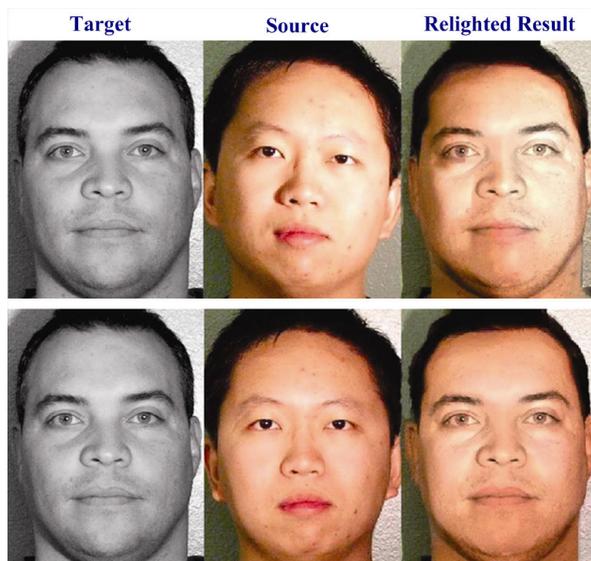


Fig. 9. Results of illumination-preserved face coloring. Left column: the target grayscale faces; middle column: the source (reference) faces with desired lighting conditions; and right column: the relighted faces.

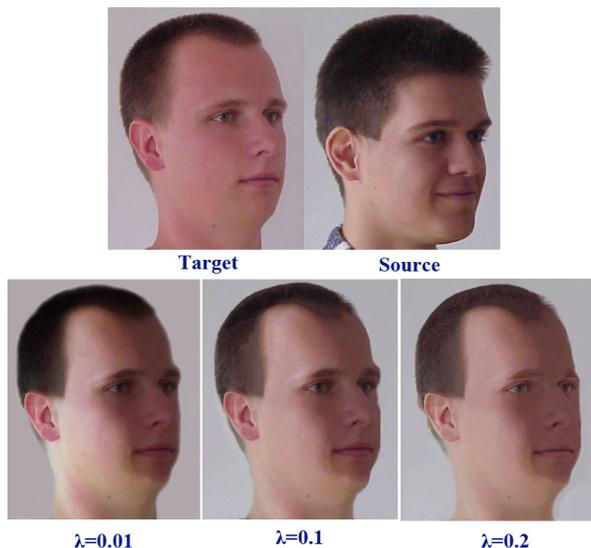


Fig. 10. An example of how to reduce visual artifacts in the relighted faces by tuning λ . The top left is the target face, the top right is the source face, the bottom row displays the relighted results when $\lambda = 0.01, 0.1, \text{ and } 0.5$, respectively.

VIII. CONCLUSIONS

We propose a novel technique for image-based face illumination transferring. Without knowing any prior lighting conditions or 3D face geometries of two 2D face images (as input), our technique is capable of automatically transferring illumination between them. Compared with the previous techniques [3]–[5], [7]–[9], our approach has the following advantages: (i) our approach does not require any pre-recorded reflectance database. In our approach, two 2D images (one is a 2D source face image with a desired lighting condition, and a target 2D face image) are the sole input. (ii) Our approach is robust and versatile. As shown in the Section VII, our approach can be used for transferring illumination across different human faces (front-view and side-view) and whether or not any of them is in grayscale or color.

The major limitations of current work include: (1) the success of this approach largely relies on the assumption that both the target and the source (reference) faces should have similar complexions, which is not always valid in some face illumination transferring applications. (2) The 2D image-based warping is required as a pre-processing

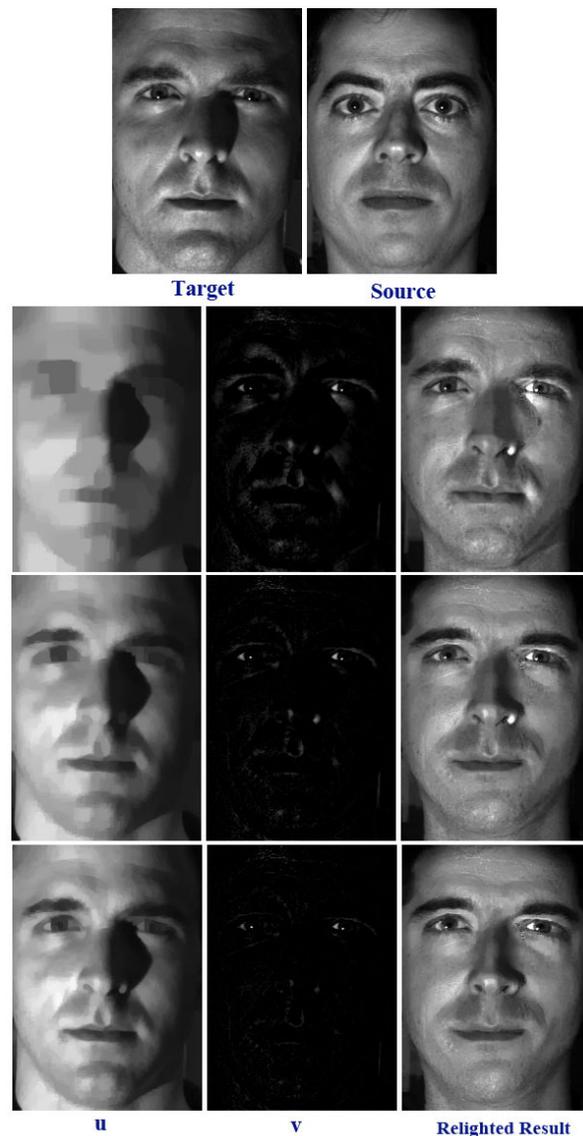


Fig. 11. The second example of how to reduce visual artifacts in the relighted faces. The first row shows two input face images (source and target). The other three rows (from top to bottom) are u component, v component, and the relighted faces when $\lambda = 0.2, 0.5,$ and $0.8,$ respectively.

step, and no 3D face geometry information has been exploited in this approach. Therefore, if there exist significant 3D face geometry differences between the source and the target face, the visual realism of the relighted face could be affected. In the future, we plan to combine automatic facial feature tracking algorithms [16], [17] to automate the face illumination transferring without the loss of perceivable visual quality.

ACKNOWLEDGMENTS

The authors would like to thank Baez Jose, Yun Chang, Qin Gu, Tanasai Sucontphunt at the Computer Graphics and Interactive Media (CGIM) Lab at the University of Houston for their data acquisition help. This research was supported by a University of Houston faculty startup fund (Dr. Zhigang Deng). W. Yin was supported by an internal faculty grant from the Dean of Engineering of Rice University.

REFERENCES

- [1] M. Nikolova, "A variational approach to remove outliers and impulse noise," *Journal of Mathematical Imaging and Vision*, vol. 20, no. 1-2, pp. 99–120, 2004.

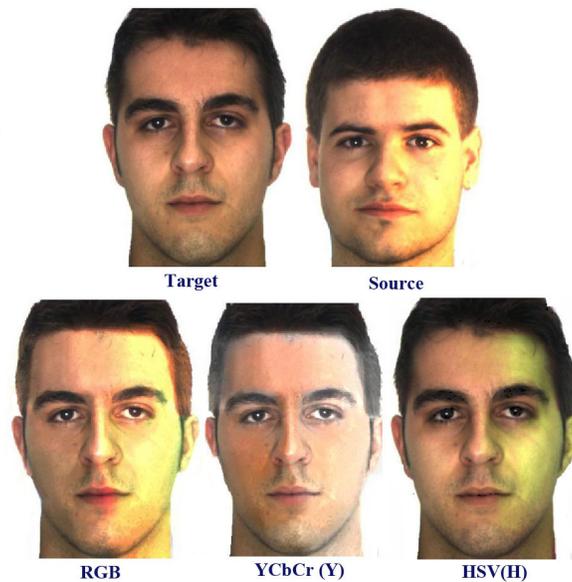


Fig. 12. An example of illumination transferring in the HSV and YCbCr color spaces. In this figure, the illumination transferring was only performed in the primary color channel (the H channel in the HSV space or the Y channel in the YCbCr space).

- [2] W. Yin, D. Goldfarb, and S. Osher, "Total variation based image cartoon-texture decomposition," *SIAM Journal on Multiscale Modeling and Simulation*, vol. 6, no. 1, pp. 190–211, 2006.
- [3] A. Stoschek, "image-based re-rendering of faces for continuous pose and illumination directions," in *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 01, 2000, pp. 582–587.
- [4] Z. Wen, Z. Liu, and T. S. Huang, "Face relighting with radiance environment maps," in *Computer Vision and Pattern Recognition, Proceedings*, 2003, pp. 157–165.
- [5] P. Peers, N. Tamura, W. Matusik, and P. Debevec, "Post-production facial performance relighting using reflectance transfer," *ACM Trans. Graph.*, vol. 26, no. 3, p. 52, 2007.
- [6] G. H. Z. W. Z. Z. Yang Wang, Zicheng Liu and D. Samaras, "Face re-lighting from a single image under harsh lighting conditions," in *Proceedings of IEEE International Conference on Computer Vision and Pattern Recognition 2007*, June 2007, pp. 1–8.
- [7] A. Shashua and T. R. Raviv, "The quotient image: Class-based recognition and synthesis under varying illumination conditions," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 02, pp. 262–265, 1999.
- [8] Z. Liu, Y. Shan, and Z. Zhang, "Expressive expression mapping with ratio images," in *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. New York, NY, USA: ACM Press, 2001, pp. 271–276.
- [9] S. R. Marschner and D. P. Greenberg, "Inverse lighting for photography," in *Proceedings of the Fifth Color Imaging Conference, Society for Imaging Science and Technology*, 1997, pp. 262–265.
- [10] T. Chen, W. Yin, X. S. Zhou, D. Comaniciu, and T. Huang, "Total variation models for variable lighting face recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 9, pp. 1519–1524, 2006.
- [11] J. Noh, D. Fidaleo, and U. Neumann, "Animated deformations with radial basis functions," in *VRST '00: Proceedings of the ACM symposium on Virtual reality software and technology*. New York, NY, USA: ACM Press, 2000, pp. 166–174.
- [12] J. P. Lewis, M. Cordner, and N. Fong, "Pose space deformation: A unified approach to shape interpolation and skeleton-driven deformation," in *Proc. of ACM SIGGRAPH'2000*, 2000, pp. 165–172.
- [13] J. Y. Noh and U. Neumann, "Expression cloning," in *Proc. of ACM SIGGRAPH'01*, 2001, pp. 277–288.
- [14] Z. Deng, P. Y. Chiang, P. Fox, and U. Neumann, "Animating blendshape faces by cross-mapping motion capture data," in *Proc. of ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, March 2006, pp. 43–48.
- [15] N. Arad, N. Dyn, D. Reifeld, and Y. Yeshurun, "Image warping by radial basis functions: Application to facial expressions," *Computer Vision, Graphics, and Image Processing. Graphical Models and Image Processing*, vol. 56, no. 2, pp. 161–172, 1994.
- [16] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active shape models—their training and application," *Comput. Vis. Image Underst.*, vol. 61, no. 1, pp. 38–59, 1995.
- [17] Y. Zhou, L. Gu, and H.-J. Zhang, "Bayesian tangent shape model: estimating shape and pose parameters via bayesian inference," in *Proceedings of IEEE Computer Society Conference*, vol. 1, 2003, pp. 109–116.
- [18] D. Goldfarb and W. Yin, "Parametric maximum flow algorithms for fast total variation minimization," *Rice University CAAM Technical Report TR07-09*, 2007.
- [19] 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. Intelligence*, vol. 23, no. 6, pp. 643–660, 2001.
- [20] "http://www.lrv.fri.uni-lj.si/," July 2007.
- [21] A. M. Martinez and R. Benavente, "The AR face database," *CVC Technical Report #24*, 1998.