Single-Image Reflectance Estimation for Relighting by Iterative Soft Grouping

Yuanzhen Li[1,2] Stephen Lin[1] Sing Bing Kang[3] Hanqing Lu[2] Heung-Yeung Shum[1]

[1]Microsoft Research, Asia [2]Chinese Academy of Sciences [3]Microsoft Research

Abstract

Reflectance values for image-based relighting are often estimated from grouped pixels with similar reflectance, but such groupings are difficult to compute with certainty for sparse image data. To address this problem, we propose an iterative method that aggregates BRDF data in a single image with known geometry and lighting by soft grouping, where pixels contribute to one another's estimate according to their degree of reflectance similarity. Estimation of specular reflectance is further improved by albedo-independent soft grouping of pixels based on shape continuity. With recovered reflectances, we demonstrate realistic relighting for synthetic and real scenes, including surfaces with spatially-varying reflectance.

1. Introduction

Photographs have commonly been used in computer graphics for realistic rendering of scenes under various illumination conditions. Previous approaches to relighting scenes include capturing a set of images under densely sampled lighting directions [2], or estimating parameters of reflectance models, which can be used for rendering with arbitrary lighting [6][9][10].

To avoid the need for large sets of images and multiple illumination conditions, some previous methods process bidirectional reflectance distribution function (BRDF) data aggregated from predefined groupings of pixels with similar reflectance [11][8][1].

While these methods assume a prior reflectance grouping, formation of these groups for reflectance estimation is a difficult task, especially when local reflectance variations often exist. To address this problem, Lensch et al. [4] present a method for clustering reflectances by iterative splitting and refitting of reflectance models to a user-specified number of materials. Spatially-varying reflectances are preserved by expressing the BRDF of each pixel in terms of basis BRDFs of its cluster. Nishino et al. [7] assemble reflectance data from point correspondences among multiple views with fixed lighting to estimate reflectance parameters as well as the illumination environment.

In our work, the goal is also to aggregate BRDF information for reflectance estimation, but with substantially reduced data. For broader applicability, our approach takes as input only a single image, and with such limited data, we cannot perform clustering as in [4] where 20-25 images are captured. In fact, a single image is insufficient for estimating the reflectance of independent pixels, so some form of data aggregation becomes necessary.

Determination of pixel groupings is, however, a challenging problem. With just a single BRDF sample for each pixel, it cannot be known with full certainty whether two pixels share the same reflectance. Furthermore, the presence of spatially-varying BRDFs will complicate the partitioning of pixels. To deal with this uncertainty, we propose a soft reflectance grouping where pixels contribute in varying degrees to one another's reflectance estimate. In computing the reflectance of a given pixel, the BRDF data of other neighboring pixels are each weighted by their reflectance similarity to the examined pixel, so that pixels which are more likely of the same reflectance are more strongly grouped while less similar pixels have relatively little impact on the estimation. By computing these soft groupings separately for each pixel, spatially-varying reflectance is also modelled in this framework.

Because of initially imprecise soft groupings, we progressively refine reflectance estimates by iterating the process. In successive iterations, updated soft groupings produce improved reflectance values, which in turn leads to more accurate comparisons of pixel similarity and better soft groupings. To further improve results, we make more complete use of BRDF information by partially grouping pixels that have similar reflectance except for albedo. On a continuous surface, it is typical for reflectance to vary only in albedo, so we take advantage of this characteristic to provide more data for estimation of non-albedo reflectance parameters. Under this scheme, we have been able to obtain reasonable reflectance estimates which have been effectively used for relighting scenes.

2. Reflectance estimation

Our algorithm for reflectance estimation takes as input a single image, geometry that can be obtained by range scans or other means, and the light source position. The reflectance models we use are a Lambertian model for diffuse reflectance and a single isotropic lobe from the Lafortune model [3] to represent other reflectance effects, such as specular reflection. In terms of light direction L, surface normal N and viewing direction V, reflectance at a point x is formulated as

$$I(x) = \rho(x)N(x)\cdot L+[c_1(L\cdot V)+c_2(N(x)\cdot L)(N(x)\cdot V)]^n \tag{1}$$

where the albedo ρ and the Lafortune coefficients c_1, c_2 and n are the four reflectance parameters we aim to recover for each image pixel.

2.1. Neighborhood support

Since estimating the reflectance of a single pixel from a single observation is an underconstrained problem, it is necessary to gather additional data to compute a feasible estimate. This data can be derived from neighboring pixels with similar reflectance, but uncertainty also exists on which pixels indeed share reflectance characteristics.

To account for this uncertainty, we propose a method with varying local support from neighboring pixels. Pixels with reflectance that seem similar to the examined pixel are more likely to share the same reflectance, and therefore their data should be emphasized in the estimation process. The less similar a pixel is to the examined pixel, the less likely it is of the same reflectance, and consequently it should have less influence on parameter estimation.

We incorporate two forms of neighborhood support in our estimation framework: reflectance similarity and shape continuity. Our measure of reflectance similarity for pixel y with respect to pixel x is based on a weighted sum of reflectance parameter differences:

$$\begin{aligned} r_x(y) &= \tfrac{1}{1+\|R(y)-R(x)\|} \\ R &= [\alpha_1\rho,\ \alpha_2 c_1,\ \alpha_3 c_2,\ \alpha_4 n]^T. \end{aligned} \tag{2}$$

where R is a vector of weighted reflectance parameters. The initial reflectance similarities are computed with respect to only albedo $\rho_k(x) = \frac{I_k(x)}{N(x)\cdot L}$ where $k = R, G, B$ and $[\alpha_1\ \alpha_2\ \alpha_3\ \alpha_4] = [1\ 0\ 0\ 0]$. After the initial iteration where the other reflectance parameters become available, we compute the reflectance similarity using all the parameters.

The reflectance for image regions shadowed from the primary light source cannot be independently estimated without at least determining the principal indirect illumination sources of each point. Because of the difficulty in this, the shadowed areas are computed from the known geometry and the primary light position, and are excluded from the reflectance estimation process. Reflectance values for these regions must nonetheless be determined for relighting purposes, so we simply associate these shadowed pixels with image pixels for which a reflectance is computed. A more refined approach to shadow generation that accounts for global illumination is presented in [5], which computes radiosity exchanges in a perfectly diffuse scene.

Shape continuity is also used in determining neighborhood support because many continuous surfaces are made of the same material, and differences in reflectance on a surface mainly result from albedo texture. We take advantage of this property to gain additional data for estimating non-albedo reflectance parameters. Our shape continuity for pixel y with respect to pixel x penalizes sharp changes in depth and surface normals along a straight-line path P between the two points:

$$s_x(y) = \textstyle\sum_{k=1}^{|P|-1}[\beta_1 \frac{1}{1+|d(P(k+1))-d(P(k))|}$$

$$+\beta_2(N(P(k+1))\cdot N(P(k))+1)^2] + r'_x(y)$$

where P is an array of pixel coordinates, d is the depth and N is the unit surface normal. β_1, β_2 are empirical weights that tradeoff depth and surface normal factors. The term $r'_x(y)$ represents the difference in specular reflectance parameters between x and y. Although r' is unrelated to shape continuity, we include this term to reduce the likelihood of grouping pixels from different materials that are present on the same continuous surface:

$$r'_x(y) = \tfrac{1}{1+\|R'(y)-R'(x)\|}$$

$$R' = [\alpha_2 c_1,\ \alpha_3 c_2,\ \alpha_4 n]^T.$$

2.2. Iterative Reflectance Estimation

The measures of reflectance similarity and shape continuity are used to drive an iterative procedure for refining reflectance estimates. The iterations are performed in two steps. In the first step, we do a soft grouping of pixels weighted by reflectance similarity. For each pixel x, its new reflectance values are computed according to

$$\arg\min_{\rho,c_1,c_2,n} \sum_{x'\in W_x} r_x(x')[I(x') - \hat{I}(x')]^2 \tag{3}$$

where W_x is a 25x25 window centered around x not including shadowed pixels, $\hat{I}$ is computed from (1), and r is from (2) with $[\alpha_1\ \alpha_2\ \alpha_3\ \alpha_4] = [1\ 0\ 0\ 0]$ for the first iteration. This equation is computed using the Levenberg-Marquardt minimization algorithm with multiple initial seeds.

In the second step, the non-albedo reflectance parameters are re-estimated with a soft grouping weighted by shape continuity. Similar to the first step, the reflectance for x is

$$\arg\min_{c_1,c_2,n} \sum_{x' \in W_x} s_x(x')[I(x') - \hat{I}(x')]^2 \qquad (4)$$

where the albedo of $\hat{I}$ is taken to be its estimated value in (3). These steps are repeated until the change in estimated reflectance values falls below a threshold t:

$$\frac{1}{|I'|} \sum_{x \in I'} [R_{new}(x) - R_{old}(x)]^2 < t$$

where I' is the set of unshadowed image pixels.

With the final reflectance estimates and scene geometry, the image can clearly be relit for arbitrary illumination conditions. Shadows are cast according to the geometry and light positions, and pixel intensities within the cast shadows result from global illumination computations [11, 5].

3. Results

We applied our method to the synthetic scene in Figure 1, which contains some challenging features such as complex texture on the container and table, and different materials that form a smoothly-shaped pipe. In our results, the scene texture is handled properly. Although specular reflection exists on only a few areas of the container in the original image, it is accurately rendered on other parts of the surface for different lighting. In particular, it can be noticed on the container that although golden-colored areas in the original image are entirely diffuse, they exhibit correct specular characteristics for other illumination conditions, since our shape-continuity soft grouping method effectively transferred reflectance properties of the material across albedos. An effect like this could not happen without albedo-independent grouping. Different materials that together form one surface, such as the pipe, are not incorrectly soft grouped because they exhibit different reflectance in the original image. The iterative refinement of reflectance estimates and its effect on relighting is illustrated in Figure 2.

We also process a real image of a ceramic cat (with range data, courtesy of Ko Nishino) that is shown in Figure 3. From the original (leftmost) image, the computed diffuse reflectance is exhibited in the second image, rendered with the same illumination as the original. The albedo values that describe the diffuse reflectance appear to fit the ceramic cat. The last two images display relighting results for two different illumination directions. The reflectance appearance in these relighting results seem reasonable, though some specular reflections are slightly broad, which results from the precision level of the geometry. Accuracy of the geometry is an important factor in obtaining realistic results.

4. Conclusion

In this paper, we have presented a method for aggregating BRDF data in a single image for estimation of reflectance parameters and subsequent relighting. To account for uncertainty in grouping pixels with limited information, our algorithm forms soft groups with variable support from neighboring pixels. Albedo-independent pixel grouping is also introduced for improving estimates of specular parameters. By computing these soft groups on a per pixel basis, spatial variations of reflectance can be captured within our estimation framework.

References

[1] S. Boivin and A. Gagalowicz. Image-based rendering of diffuse, specular and glossy surfaces from a single image. *SIGGRAPH '01 Conference Proceedings*, pages 107–116, 2001.

[2] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. *SIGGRAPH '00 Conference Proceedings*, pages 145–156, 2000.

[3] E. P. Lafortune, S.-C. Foo, K. E. Torrance, and D. P. Greenberg. Non-linear approximation of reflectance functions. *SIGGRAPH '97 Conference Proceedings*, pages 117–126, 1997.

[4] H. P. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel. Image-based reconstruction of spatially-varying materials. *Eurographics Rendering Workshop*, pages 104–115, June 2001.

[5] C. Loscos, G. Drettakis, and L. Robert. Interactive virtual relighting of real scenes. *IEEE Transactions on Visualization and Computer Graphics*, 6(3):289–305, 2000.

[6] S. R. Marschner, S. H. Westin, E. P. Lafortune, K. E. Torrance, and D. E. Greenberg. Image-based brdf measurement including human skin. *Eurographics Rendering Workshop*, pages 131–144, June 1999.

[7] K. Nishino, Z. Zhang, and K. Ikeuchi. Determining reflectance parameters and illumination distribution from a sparse set of images for view-dependent image synthesis. *International Conference on Computer Vision*, 1:599–606, July 2001.

[8] R. Ramamoorthi and P. Hanrahan. A signal-processing framework for inverse rendering. *SIGGRAPH '01 Conference Proceedings*, pages 117–128, 2001.

[9] Y. Sato, M. D. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. *SIGGRAPH '97 Conference Proceedings*, pages 379–387, 1997.

[10] T.-T. Wong, P.-A. Heng, S.-H. Or, and W.-Y. Ng. Image-based rendering with controllable illumination. *Eurographics Rendering Workshop*, pages 13–22, June 1997.

[11] Y. Yu, P. E. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs. *SIGGRAPH '99 Conference Proceedings*, pages 215–224, 1999.

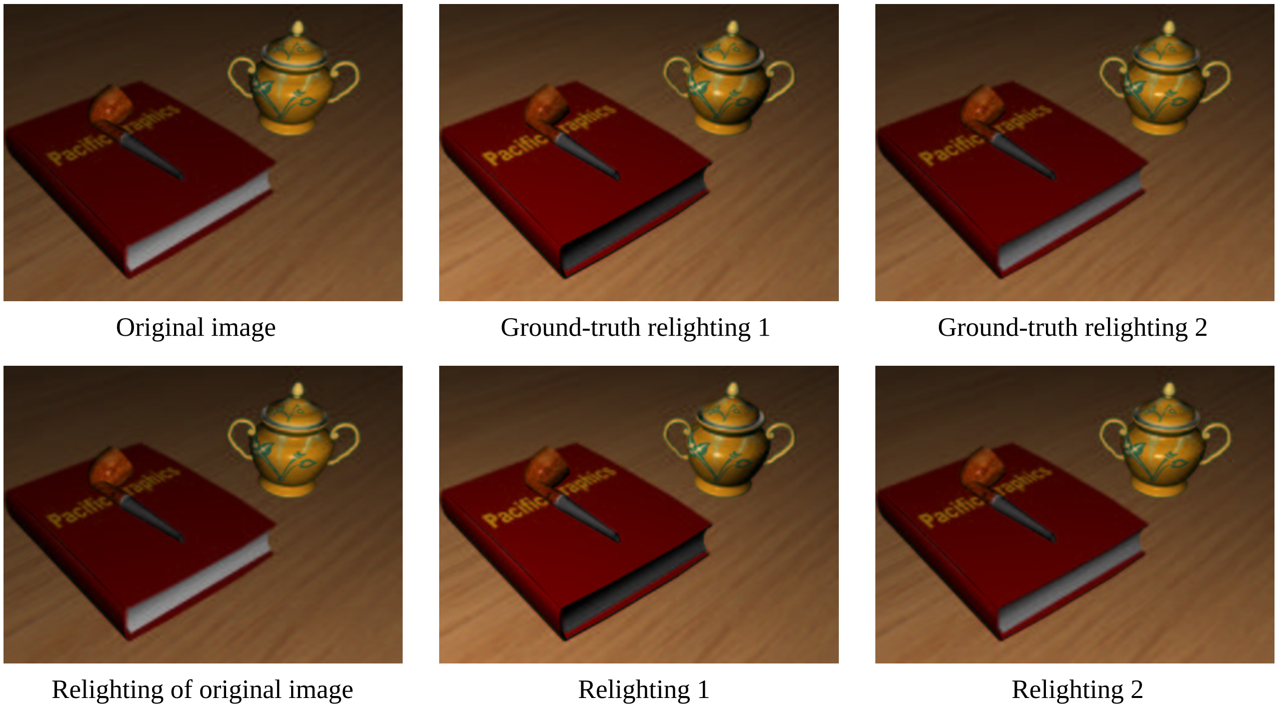

Figure 1. Synthetic images and relighting results

Intermediate relighting 1 | Intermediate relighting 2 | Final relighting | Ground-truth relighting

Figure 2. Effect of iterations on relighting

Original image | Diffuse image | Relighting 1 | Relighting 2

Figure 3. Real image and relighting results