

Cosine Lobe Based Relighting from Gradient Illumination Photographs

Graham Fyffe, Paul Debevec

University of Southern California Institute for Creative Technologies
12015 Waterfront Drive, Playa Vista, CA 90094, USA
Email: {fyffe, debevec}@ict.usc.edu FAX: +1-310-574-5725

Abstract

We present an image-based method for relighting a scene by analytically fitting cosine lobes to the reflectance function at each pixel, based on gradient illumination photographs. Realistic relighting results for many materials are obtained using a single per-pixel cosine lobe obtained from just two color photographs: one under uniform white illumination and the other under colored gradient illumination. For materials with wavelength-dependent scattering, a better fit can be obtained using independent cosine lobes for the red, green, and blue channels, obtained from three achromatic gradient illumination conditions instead of the colored gradient condition. We explore two cosine lobe reflectance functions, both of which allow an analytic fit to the gradient conditions. One is non-zero over half the sphere of lighting directions, which works well for diffuse and specular materials, but fails for materials with broader scattering such as fur. The other is non-zero everywhere, which works well for broadly scattering materials and still produces visually plausible results for diffuse and specular materials. We also perform an approximate diffuse/specular separa-

tion of the reflectance, and estimate scene geometry from the recovered photometric normals to produce hard shadows cast by the geometry, while still reconstructing the input photographs exactly.

1 Introduction and Related Work

Image-based relighting is a powerful technique for synthesizing images of a scene under novel illumination conditions, based on a set of input photographs. In its most basic form, images of a scene are acquired [Hae92] or rendered [NSD94] under a set of basis lighting conditions. Then, a relit version of the scene can be produced by taking linear combinations of the basis lighting conditions, akin to compositing together different lighting passes of a model miniature. Debevec et al. [DHT⁺00] used a light stage device to acquire a dataset of a human face lit by a dense set of over two thousand lighting directions on the sphere, and showed that such datasets could be efficiently illuminated under novel real-world lighting conditions such as high dynamic range lighting environments through image-based relighting. Recent work in the area of precomputed radiance transfer [SKS02], [RH02], [WTL05] has shown that pre-rendering an object's reflectance under basis illumination conditions can allow for real-time relighting as it moves through interactive environments. Basis illumination datasets have also been shown to be very useful for object recognition [Ram06], including for faces.

The principal benefit to image-based relighting techniques is that complex illumination effects including spatially-varying diffuse and specular reflection, self-shadowing, mutual illumination, and subsurface scattering are all encoded within the data and thus appear accurately in the renderings, whereas traditional

Digital Peer Publishing Licence

Any party may pass on this Work by electronic means and make it available for download under the terms and conditions of the current version of the Digital Peer Publishing Licence (DPPL). The text of the licence may be accessed and retrieved via Internet at <http://www.dipp.nrw.de/>.

First presented at the European Conference on Visual Media Production (CVMP) 2009, extended and revised for JVRB

techniques would require far more advanced reflectometry and light transport simulation. The principal drawbacks are that a lot of data must be acquired and stored. This makes the techniques less practical for capturing dynamic subjects (high-speed video at thousands of frames per second is required for dynamic subjects as in [WGT⁺05]) and for memory-efficient relighting as hundreds of images can be required.




Figure 1: A reflectance function $f(\Theta)$ maps incoming radiance directions Θ at a point \mathbf{p} on a surface to the resulting radiance that is reflected along the view ray $\hat{\mathbf{v}}$ towards a camera.

Prior works have explored efficient representations of image-based relighting datasets for rendering, generally by focusing on representations for the scene's per-pixel *reflectance functions* (Fig. 1) which map incoming radiance direction on the sphere to the resulting radiance (i.e. pixel color) that is reflected towards a fixed camera position. Debevec et al. [DHT⁺00] estimated diffuse and specular albedos and normals for each pixel's reflectance function, reducing the information for each pixel from hundreds of reflectance measurements to just a few reflectance parameters. However, these parameters attempted to factor out global illumination effects, requiring these effects to be simulated later and forfeiting the elegance and realism of image-based relighting. Malzbender et al. [MGW01] fit quadratic *polynomial texture maps* (PTMs) to reflectance functions consisting of fifty lighting directions across the hemisphere. The PTMs greatly reduce the reflectance data to a compact, data-driven representation, and resulting renderings produce realistic and relatively smooth and diffuse renditions of the objects under varying illumination. However, the technique still requires a dense set of incident lighting directions to be recorded. Also, it restricted its consideration to lighting originating from the front hemisphere, which is a significant limitation for fully three-dimensional objects. Ma et al. [MHP⁺07] used

a computational illumination approach to modeling reflectance functions using a small number of incident lighting conditions. Using a spherical light stage and a set of four spherical gradient illumination conditions derived from the 0th- and 1st-order spherical harmonics, the technique directly measures the magnitude (albedo) and centroid (surface normal) of each pixel's reflectance function. These measurements are used to drive a Lambertian or Phong reflectance lobe to represent the reflectance function; using polarization difference imaging, the diffuse and specular components can be modeled independently. As with PTMs, the resulting renderings still encode most of the effects of global illumination. However, the lobe widths needed to be selected manually (either choosing a Lambertian lobe or a Phong specular exponent), as no reflectance lobe width information was derived from the measurements. Very recent work by Ghosh et al. [GCP⁺09] has estimated per-pixel reflectance lobe widths by adding 2nd-order spherical harmonic illumination conditions to the basis of [MHP⁺07], and also estimates anisotropic reflection parameters. However, this approach requires increasing the number of illumination conditions to at least seven even if reflections are assumed to be isotropic.

In this paper, we note that the four illumination conditions of [MHP⁺07] actually *overdetermine* the surface albedo (one parameter per color) and surface normal or reflected direction (two parameters per color) of each pixel's reflectance, and we show that the remaining degree of freedom holds information about the breadth of the reflectance lobe. We present two new data-driven reflectance models based on cosine lobes which can be fit analytically to measurements under the four gradient illumination conditions, one which is appropriate for opaque materials and one for semi-translucent materials. To reduce the number of lighting conditions from four to just two, we follow the suggestion of Woodham [Woo80] and encode three different illumination patterns into the red, green, and blue color channels of incident illumination using a color-enabled light stage as in [DWT⁺02]. We obtain surface normal estimates through an approximate diffuse/specular separation of the reflectance, and we further enable the simulation of high-frequency cast shadows in the renderings by estimating scene geometry from the surface normal estimates. The result is an integrated technique which allows relatively detailed reflectance acquisition and simulation from an extremely small number of incident lighting conditions.

2 Method

We photograph the scene inside a geodesic sphere of colored LED lights, or Light Stage, which can be programmed to produce gradient illumination as well as uniform white illumination, similar to [DWT⁺02]. The Light Stage makes use of 156 iColor MR lights from Color Kinetics corporation, which are individually controllable from a single computer via USB.

We explore two different cosine lobe reflectance functions, which integrate easily over the uniform and gradient illumination conditions, so that an analytic fit to the observations is obtained. *Hemispherical* cosine lobes of the form $f(\Theta) = k(\hat{\alpha} \cdot \Theta)^n$ are explored for modeling diffuse and specular materials, but are unable to represent materials with broader scattering such as fur. Alternatively, *spherical* cosine lobes of the form $f(\Theta) = k(\frac{1}{2}\hat{\alpha} \cdot \Theta + \frac{1}{2})^n$ are explored for broadly scattering materials, and are still flexible enough to offer some representation of diffuse and specular materials. In both models, $\hat{\alpha}$ refers to the axis of the lobe, while k and n are constants.

Many real-world reflectance functions are better fit by two lobes instead of one, being the so-called “diffuse” and “specular” lobes. We perform an approximate diffuse/specular separation to recover independent diffuse and specular lobes, as well as a surface normal estimate. Finally, for offline relighting we attempt to introduce high-frequency features into the reflectance functions in order to produce hard cast shadows in relit images. We first estimate the scene geometry by integrating photometric normals, and then estimate high-frequency reflectance functions based on the visibility implied by the geometry, adjusted to maintain exact reconstruction of the input photographs. These steps are described in detail in the sections that follow.

2.1 Data Acquisition

In the data acquisition stage, we make four observations o_w, o_x, o_y, o_z for each pixel and each color channel. The four observed illumination conditions defined over the unit sphere of directions $\Theta \in S^2$ with overall illumination intensity L are $L, (\frac{1}{2}\Theta_x + \frac{1}{2})L, (\frac{1}{2}\Theta_y + \frac{1}{2})L,$ and $(\frac{1}{2}\Theta_z + \frac{1}{2})L,$ corresponding to a full-on illumination condition and three orthogonal linear gradient illumination conditions. Refer to Fig. 2 for a subject photographed under these four illumination conditions. In the case of four input photographs, the four observations are trivially the input photograph




Figure 2: Cat subject photographed under the full-on and gradient illumination conditions




Figure 3: Top: Cat subject photographed under the full-on (A) and colored gradient (B) illumination conditions. Middle: The red, green, and blue components of the ratio of (B) to (A). Bottom: The middle row multiplied by (A). Note the resemblance to Fig. 2.

pixel values. If a single colored gradient is used instead of separate x, y and z gradients, we synthesize the observations o_x, o_y, o_z by multiplying the full-on image with each of the three color channels of the ra-

tio image of the colored gradient image to the full-on image, motivated by [Woo80]. This process is illustrated in Fig. 3. (Compare the images in Fig. 2 with the bottom row of Fig. 3.) Additionally, since the color primaries of the LEDs are not exactly the same as the color primaries of the camera sensors, we calibrate a color correction matrix prior to acquisition in order to reduce cross-talk between color channels.

2.2 Cosine Lobe Reflectance Functions

It is informative to first examine the general case of the integral of a hemispherical cosine lobe $k(\hat{\alpha} \cdot \Theta)^n$ times an arbitrary linear gradient $\beta \cdot \Theta + b$ (with β not necessarily unit) over the hemisphere $\Omega_+(\hat{\alpha})$ of directions on the positive side of the axis of the lobe. The integral has the closed form:

$$\int_{\Omega_+(\hat{\alpha})} k(\hat{\alpha} \cdot \Theta)^n (\beta \cdot \Theta + b) d\omega_{\Theta} = \frac{2\pi k(\hat{\alpha} \cdot \beta)}{n+2} + \frac{2\pi k b}{n+1}. \quad (1)$$

From here we trivially derive the four illumination condition observations:

$$\begin{aligned} \int_{\Omega_+(\hat{\alpha})} k(\hat{\alpha} \cdot \Theta)^n L d\omega_{\Theta} &= \frac{2\pi k L}{n+1}, \\ \int_{\Omega_+(\hat{\alpha})} k(\hat{\alpha} \cdot \Theta)^n \left(\frac{1}{2}\Theta_x + \frac{1}{2}\right) L d\omega_{\Theta} &= \frac{\pi k L \hat{\alpha}_x}{n+2} + \frac{\pi k L}{n+1}, \\ \int_{\Omega_+(\hat{\alpha})} k(\hat{\alpha} \cdot \Theta)^n \left(\frac{1}{2}\Theta_y + \frac{1}{2}\right) L d\omega_{\Theta} &= \frac{\pi k L \hat{\alpha}_y}{n+2} + \frac{\pi k L}{n+1}, \\ \int_{\Omega_+(\hat{\alpha})} k(\hat{\alpha} \cdot \Theta)^n \left(\frac{1}{2}\Theta_z + \frac{1}{2}\right) L d\omega_{\Theta} &= \frac{\pi k L \hat{\alpha}_z}{n+2} + \frac{\pi k L}{n+1}. \end{aligned}$$

Substituting the observations o_w, o_x, o_y, o_z corresponding to the four lighting conditions and requiring $\hat{\alpha}$ to be a unit vector yields a system of five equations:

$$\begin{aligned} o_w &= \frac{2\pi k L}{n+1}, \\ o_x &= \frac{\pi k L \hat{\alpha}_x}{n+2} + \frac{\pi k L}{n+1}, \\ o_y &= \frac{\pi k L \hat{\alpha}_y}{n+2} + \frac{\pi k L}{n+1}, \\ o_z &= \frac{\pi k L \hat{\alpha}_z}{n+2} + \frac{\pi k L}{n+1}, \\ \|\hat{\alpha}\| &= 1. \end{aligned}$$

These equations admit a single closed-form solution for the cosine lobe:

$$\hat{\alpha} = \frac{\alpha}{\|\alpha\|}, \quad (2a)$$

$$n = \frac{2\|\alpha\| - o_w}{o_w - \|\alpha\|}, \quad (2b)$$

$$k = \frac{o_w(n+1)}{2\pi L}, \quad (2c)$$

where $\alpha = \langle 2o_x - o_w, 2o_y - o_w, 2o_z - o_w \rangle$.

Note that the cosine lobe axis is obtained from the gradient illumination condition observations in a manner similar to the method in [MHP⁺07] for obtaining the mean spherical angle of a reflectance function. This is due to the symmetry of the cosine lobe reflectance function about its axis.

By a similar argument, a spherical cosine lobe of the form $k(\frac{1}{2}\hat{\alpha} \cdot \Theta + \frac{1}{2})^n$ has the solution:

$$\hat{\alpha} = \frac{\alpha}{\|\alpha\|}, \quad (3a)$$

$$n = \frac{2\|\alpha\|}{o_w - \|\alpha\|}, \quad (3b)$$

$$k = \frac{o_w(n+1)}{4\pi L}. \quad (3c)$$




Figure 4: Coefficients for the cat subject under the hemispherical cosine lobe model. Top, left to right: $\hat{\alpha}_x, \hat{\alpha}_y, \hat{\alpha}_z$. Bottom, left to right: k, n , zoom of n . The value of the exponent n in the figure is scaled by $\frac{1}{8}$.

We store only the coefficients of the cosine lobe reflectance model at each pixel for each color channel. Fig. 4 visualizes the coefficients for a cat subject, where the red, green, and blue channel coefficients are grouped and displayed as color images. Once the cosine lobes for each pixel and each color channel



Figure 5: Relit images under a novel point-light illumination condition. The light is coming from far right. Left: hemispherical cosine lobe model. Right: spherical cosine lobe model.

are computed, relit images can easily be synthesized for any distant illumination condition. Relit images are computed simply by sampling the cosine lobe reflectance functions for each light source in the novel illumination condition. See Fig. 5 for examples of relit images generated by this method.

2.3 Diffuse/Specular Separation

The four observations under different illumination conditions in general do not contain enough information to be separated into diffuse and specular components. However, under certain assumptions about isotropy and the breadth of the diffuse and specular lobes, a separation can indeed be obtained. Notably, other than the assumptions just mentioned, this separation does not depend on any other assumptions about the shape of the lobes, so it may find applications beyond cosine lobe reflectance models.

We presume the observations o to be a sum of diffuse and specular components d and s :

$$o = d + s.$$

For each color channel, we consider a quantity g describing the gradient illumination observations normalized by the full-on illumination observation:

$$g = \frac{1}{o_w} \langle 2o_x - o_w, 2o_y - o_w, 2o_z - o_w \rangle.$$

We likewise consider quantities d and s describing the normalized diffuse and specular components:

$$d = \frac{1}{d_w} \langle 2d_x - d_w, 2d_y - d_w, 2d_z - d_w \rangle;$$

$$s = \frac{1}{s_w} \langle 2s_x - s_w, 2s_y - s_w, 2s_z - s_w \rangle.$$

It follows that:

$$g = \frac{d_w}{d_w + s_w} d + \frac{s_w}{d_w + s_w} s,$$

and hence g lies on the line segment \overline{ds} .

If we assume that the reflectance is isotropic, and that the diffuse component d is aligned with the surface normal n , then the vector v pointing towards the viewer and the vectors g , d and s must all lie in the same plane. Hence we define a coordinate system in the plane defined by v and (mean) g , and consider only two-dimensional vectors from here on by projecting onto this plane. On this plane, the angle Θ between v and n is equal to the angle between n and the ideal reflection vector r . Fig. 6 depicts the geometric relationships between all of these vectors. Additionally, d is the intersection of n with a circle of radius $|d|$, and s is the intersection of r with a circle of radius $|s|$.




Figure 6: On the plane defined by the view vector v and the (mean) measured vector g : the red, green, and blue dots are the normalized gradient vectors g of the red, green and blue color channels; n is the surface normal; r is the ideal reflection vector; d and s are the normalized diffuse and specular components; and D and S are circles of length $|d|$ and $|s|$. If $|d|$ and $|s|$ are known, then Θ is uniquely determined as the angle that yields the best fit for the line segment \overline{ds} to the vectors g .

If $|d|$ and $|s|$ are known, but n and r are unknown, we can recover n and r (and hence d and s) by finding the angle Θ that yields the best fit for the line segment \overline{ds} to the vectors g of each color component. For example, by Eq. (2) it can be shown that $|d| = \frac{2}{3}$ for

ideal Lambertian (diffuse) reflectance, and $|s| = 1$ for ideal specular (perfect mirror) reflectance. With $|d|$ and $|s|$ fixed, we find just one Θ to explain all three color channels. As the behavior of the line segment \overline{ds} is non-linear, no closed form for Θ could be obtained and hence we find it using a simple line search.

After Θ is found, we project the vectors g of each color channel onto the line segment \overline{ds} to compute the (per-color-channel) quantity $\frac{s_w}{d_w+s_w}$, which allows us to recover the un-normalized diffuse and specular components d and s . However, since the data may deviate from our assumptions in several respects, the vectors g may not lie exactly on the line segment \overline{ds} , and therefore the reconstruction of the data using d and s may not be exact. Therefore we let $d = g - s$, which has the effect of lumping the residuals of the reconstruction into the diffuse component, which seems acceptable as occlusion and interreflection already pollute the diffuse component to some degree.

Finally, with the recovered diffuse and specular components d and s for each illumination condition, we fit independent diffuse cosine lobes and specular cosine lobes. By this construction, the diffuse and specular cosine lobes reconstruct the input observations exactly. See Fig. 7 for an example of diffuse/specular separation under the Lambertian / ideal specular assumption.




Figure 7: Diffuse / specular separation of the cat. Top, left to right: original full-on illumination photograph, recovered normals. Bottom, left to right: recovered diffuse component, recovered specular component.

2.4 Estimating Geometry and Shadows

The estimated cosine lobe reflectance function $f(\Theta)$ has only low-frequency features. Thus, while the low-frequency-illumination input photographs are reconstructed exactly, hard cast shadows cannot be synthesized for novel illumination conditions having high-frequency features. To remedy this, we estimate scene geometry by integrating the surface normals obtained in the diffuse/specular separation phase. Integration proceeds similar to [HB86]. See Fig. 8 and 9 for examples of reconstructed geometry. After the scene geometry is estimated, we compute an adjusted reflectance function $f'(\Theta)$ that exhibits hard shadows




Figure 8: Reconstructed geometry from photometric normals. Top: Two views of the cat subject. Bottom: Two views of the duck subject.



Figure 9: Reconstructed geometry for the cat subject, with normals derived from two photos instead of four.

cast by the geometry. First, consider a reflectance function $f^*(\Theta)$ adjusted simply by occluding any light directions that hit scene geometry:

$$f^*(\Theta) = f(\Theta) \text{ visible}(\Theta), \quad (4)$$

where $\text{visible}(\Theta)$ is 0 if the direction Θ is occluded by geometry or outside of the domain of $f(\Theta)$, and 1 otherwise. To reduce shadow aliasing, we allow $\text{visible}(\Theta)$ to fall off gradually from 1 to 0 near the surface of the geometry, so that a ray that just grazes the geometry will be attenuated, but still continue through the scene. In general, relighting using $f^*(\Theta)$ will not reconstruct the input photographs exactly, since the low-frequency component of the reflectance function will be altered by the visibility masking. To correct this, we introduce a scaling factor c to restore the overall reflectance magnitude, and introduce a low-frequency “ambient” term $\rho \cdot \Theta$ to restore the original low-frequency component of the reflectance function:

$$f'(\Theta) = cf^*(\Theta) + \rho \cdot \Theta. \quad (5)$$

Let $o_w^*, o_x^*, o_y^*, o_z^*$ represent synthetic observations using $f^*(\Theta)$, computed by numerical integration:

$$\begin{aligned} o_w^* &= \int_{S^2} f^*(\Theta) L d\omega_\Theta, \\ o_x^* &= \int_{S^2} f^*(\Theta) \left(\frac{1}{2}\Theta_x + \frac{1}{2}\right) L d\omega_\Theta, \\ o_y^* &= \int_{S^2} f^*(\Theta) \left(\frac{1}{2}\Theta_y + \frac{1}{2}\right) L d\omega_\Theta, \\ o_z^* &= \int_{S^2} f^*(\Theta) \left(\frac{1}{2}\Theta_z + \frac{1}{2}\right) L d\omega_\Theta. \end{aligned}$$

Then we correct the observations using c, ρ :

$$\begin{aligned} o_w &= co_w^* + \int_{S^2} (\rho \cdot \Theta) L d\omega_\Theta, \\ o_x &= co_x^* + \int_{S^2} (\rho \cdot \Theta) \left(\frac{1}{2}\Theta_x + \frac{1}{2}\right) L d\omega_\Theta, \\ o_y &= co_y^* + \int_{S^2} (\rho \cdot \Theta) \left(\frac{1}{2}\Theta_y + \frac{1}{2}\right) L d\omega_\Theta, \\ o_z &= co_z^* + \int_{S^2} (\rho \cdot \Theta) \left(\frac{1}{2}\Theta_z + \frac{1}{2}\right) L d\omega_\Theta, \end{aligned}$$

which admit one closed-form solution:

$$c = \frac{o_w}{o_w^*}, \quad (6a)$$

$$\rho = \frac{3}{2\pi L} \langle o_x - co_x^*, o_y - co_y^*, o_z - co_z^* \rangle. \quad (6b)$$

The plausibility of the low-frequency term $\rho \cdot \Theta$ can be justified as follows: If a ray cast in the direction Θ hits scene geometry, then the only way Θ can contribute to the resulting radiance is by indirect bounce light, as illustrated in Fig. 10. Indirect bounce light is often low-frequency in nature, so it seems appropriate to model it by a low-frequency term.




Figure 10: If a ray cast in the direction Θ hits scene geometry, then the only way Θ can contribute to the resulting radiance is by indirect bounce light.

Relighting now proceeds by sampling $f'(\Theta)$ for each light. To avoid storing the entire function $f'(\Theta)$ at each pixel, we instead store the scene geometry as a depth map, the original estimated reflectance function $f(\Theta)$, and the correction coefficients c and ρ , and use ray casting to evaluate $\text{visible}(\Theta)$ as needed. Errors in the scene geometry, especially in the background plane, can introduce unwanted shadows into the relit images. To reduce this problem at the expense of some manual intervention, we introduce a user-supplied binary mask to prevent the background elements from occluding the foreground elements. Fig. 11 shows such a mask, and compares the results obtained with and without using the mask.



Figure 11: A user-supplied visibility mask (center) reduces unwanted shadows. Left column: Relit images under two point-light illumination conditions, no mask supplied. Right column: Mask supplied.




Figure 12: Relit images of cat subject under novel point light illumination conditions. A) Linear basis using four photographs. B) Hemispherical cosine lobe using two photographs. C) Hemispherical cosine lobe using four photographs. D) Diffuse and specular hemispherical cosine lobes using four photographs. E) Diffuse and specular hemispherical cosine lobes with hard cast shadows, using four photographs and user-supplied visibility mask. F) Ground truth, with inset gray balls indicating illumination direction.




Figure 13: Relit images of duck subject under novel point light illumination conditions. A) Linear basis using four photographs. B) Spherical cosine lobes using two photographs. C) Spherical cosine lobes using four photographs. D) Ground truth.

3 Results and Discussion

Refer to Fig. 12 and 13 for ground-truth comparisons with several relighting methods on the cat and duck subject, respectively. For both subjects, a linear basis relighting method is included as a baseline, which attempts to synthesize novel illumination conditions as a linear combination of the four observed gradient conditions. Note that all relighting methods tested here, including the linear basis method, reconstruct the set of input photographs exactly.

The cat subject relit with the linear basis reference method appears flat and does not convey the shape of the cat well. The two-photograph hemispherical cosine lobe method has noticeably lower quality than the four-photograph hemispherical cosine lobe method, but could be suitable for applications with limited capture budget such as video relighting. Separating the reflectance into diffuse and specular lobes significantly improves the result for the cat: the behavior of the materials under back-lighting is plausible, and specular highlights are more pronounced. Notably, the dif-

fuse/specular separation visualized in Fig. 7 largely agrees with the materials of the actual subject. For example, the golden body of the cat is a specular material with dependency on angle, in this case a gold paint. The red collar, ears, and nose have a dominant diffuse red component with a small specular component in all channels, in this case a semi-gloss red paint. The green bib has a dominant diffuse green component, in this case a very soft diffuse green paint with a faint gloss. Most notable is the visual softness of the materials in the diffuse component. Almost without exception, the specular highlights in the original photograph are separated into the specular component, leaving a smooth diffuse component. The most visible examples include the flecks of glitter on the green bib, and the highlights in the eyes. The reconstructed geometry clearly suffers from depth ambiguities in the normal integration, but in a relighting scheme that does not attempt to alter the viewpoint, the hard cast shadows and self shadowing appear to be convincing in most cases. However, the requirement of a user-supplied visibility mask to avoid unwanted shadows in some of the lighting con-

ditions is a notable limitation of the method, and future work should attempt to generate such masks automatically, or to improve the geometry estimate to remove the need for such masks.

The duck subject relit with the linear basis reference method performs better than on the cat, but does exhibit some disturbing shadowing artifacts. The two-photograph and four-photograph spherical cosine lobe reconstructions of the duck both reproduce the scattering of the fur well in most places, and still provide a good approximation of the other materials in the scene, including the specular necktie, but some lack of self shadowing is evident. Hard cast shadows for the duck fail altogether, as shown in Fig. 14, because the visibility masking assumption does not hold for highly scattering materials. Better self-shadowing across a wider range of materials would be an improvement sought after in future work, perhaps using a more complete estimate of visibility based on the reflectance function estimates of all pixels.



Figure 14: The visibility masking assumption does not hold for highly scattering materials such as the duck subject, resulting in an inappropriately hard shadow.

4 Conclusion

We presented a method for estimating the reflectance of a scene for the purpose of relighting, using just two to four photographs of the scene under gradient illumination. The gradient illumination photographs provide enough information to fit a cosine-lobe reflectance function, including an exponent. The cosine lobes produce plausible results for materials with reflectance dominated by a single lobe, regardless of whether the lobe results from reflection, scattering, or retro-reflection. Improved results including self-shadowing are obtained for materials having a Lambertian diffuse component and highly specular component via an approximate diffuse/specular separation. Ground truth comparisons reveal several limitations with our method, mostly stemming from violations of the assumptions made in our reflectance model. In fu-

ture work we would aim to broaden the applicability of our method to a wider variety of real-world materials, to enable geometry estimation and self-shadowing for materials that fall outside of the Lambertian plus specular reflectance model.

5 Acknowledgments

The authors wish to thank Jay Busch and Cyrus Wilson for helping prepare the manuscript. We would also like to thank Monica Nicholson, Bill Swartout, and Randall Hill for their support and assistance with this work.

References

- [DHT⁺00] Paul Debevec, Tim Hawkins, Chris Tchou, Haarm-Pieter Duiker, Westley Sarokin, and Mark Sagar, *Acquiring the Reflectance Field of a Human Face*, Proc. ACM SIGGRAPH 2000, Computer Graphics Proceedings, Annual Conference Series, 2000, pp. 145–156, ISBN 1581132085.
- [DWT⁺02] Paul Debevec, Andreas Wenger, Chris Tchou, Andrew Gardner, Jamie Waese, and Tim Hawkins, *A Lighting Reproduction Approach to Live-Action Compositing*, ACM Transactions on Graphics **21** (2002), no. 3, 547–556, ISSN 0730-0301.
- [GCP⁺09] Abhijeet Ghosh, Tongbo Chen, Pieter Peers, Cyrus A. Wilson, and Paul Debevec, *Estimating Specular Roughness and Anisotropy from Second Order Spherical Gradient Illumination*, Computer Graphics Forum, vol. 28, 2009, pp. 1161–1170.
- [Hae92] Paul Haeberli, *Synthetic Lighting for Photography*, <http://www.graficaobscura.com/synth/> (Last visited September 29th, 2010), January 1992.
- [HB86] Berthold K. P. Horn and Michael J. Brooks, *The variational approach to shape from shading*, Computer Vision, Graphics and Image Processing **33** (1986), no. 2, 174–208, ISSN 0734-189X.

- [MGW01] Tom Malzbender, Dan Gelb, and Hans Wolters, *Polynomial Texture Maps*, Proc. SIGGRAPH 2001, 2001, pp. 519–528, ISBN 1-58113-292-1.
- [MHP⁺07] Wan-Chun Ma, Tim Hawkins, Pieter Peers, Charles-Felix Chabert, Malte Weiss, and Paul Debevec, *Rapid Acquisition of Specular and Diffuse Normal Maps from Polarized Spherical Gradient Illumination*, Rendering Techniques 2007: 18th Eurographics Symposium on Rendering, 2007, pp. 183–194.
- [NSD94] Jeffrey S. Nimeroff, Eero Simoncelli, and Julie Dorsey, *Efficient Re-rendering of Naturally Illuminated Environments*, Fifth Eurographics Workshop on Rendering, 1994, pp. 359–373.
- [Ram06] Ravi Ramamoorthi, *Modeling Illumination Variation with Spherical Harmonics*, Face Processing: Advanced Modeling Methods (Wenyi Zhao and Rama Chellapa, eds.), Academic Press, 2006, pp. 385–424, ISBN 0120884526.
- [RH02] Ravi Ramamoorthi and Pat Hanrahan, *Frequency Space Environment Map Rendering*, ACM Transactions on Graphics **21** (2002), no. 3, 517–526, ISSN 0730-0301.
- [SKS02] Peter-Pike Sloan, Jan Kautz, and John Snyder, *Precomputed Radiance Transfer for Real-Time Rendering in Dynamic, Low-Frequency Lighting Environments*, ACM Transactions on Graphics **21** (2002), no. 3, 527–536, ISSN 0730-0301.
- [WGT⁺05] Andreas Wenger, Andrew Gardner, Chris Tchou, Jonas Unger, Tim Hawkins, and Paul Debevec, *Performance relighting and reflectance transformation with time-multiplexed illumination*, ACM Transactions on Graphics **24** (2005), no. 3, 756–764, ISSN 0730-0301.
- [Woo80] Robert J. Woodham, *Photometric Method for Determining Surface Orientation from Multiple Images*, Optical Engineering **19** (1980), no. 1, 139–144, ISSN 0091-3286.
- [WTL05] Rui Wang, John Tran, and David Luebke, *All-frequency interactive relighting of translucent objects with single and multiple scattering*, ACM Transactions on Graphics **24** (2005), no. 3, 1202–1207, ISSN 0730-0301.

Citation
Graham Fyffe and Paul Debevec, <i>Cosine Lobe Based Relighting from Gradient Illumination Photographs</i> , Journal of Virtual Reality and Broadcasting, Volume 9(2012), no. 2, February 2012, urn:nbn:de:0009-6-32649, ISSN 1860-2037.