

Cosine Lobe Based Relighting from Gradient Illumination Photographs

Graham Fyffe Cyrus A. Wilson Paul Debevec
University of Southern California Institute for Creative Technologies
13274 Fiji Way, 5th Floor, Marina del Rey, CA 90292, USA
Email: {fyffe,cwilson,debevec}@ict.usc.edu FAX: +1-310-577-9140

Abstract—We present an image-based method for relighting a scene by analytically fitting a cosine lobe 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 monochromatic 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. Additionally, we estimate scene geometry from the photometric normals to produce hard shadows cast by the geometry, while still reconstructing the input photographs exactly.

I. 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 (Haeberli [1]) or rendered (Nimeroff et al. [2]) 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. [3] 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 (Sloan et al. [4], Ramamoorthi et al. [5], and Wang et al. [6]) 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 (Ramamoorthi [7]), 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 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 Wenger et al. [8]) and for memory-efficient relighting as hundreds of images can be required.

Some work has explored efficient representations of Image-Based Relighting datasets for rendering, generally by focusing on efficient 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 camera. Debevec et al. [3] 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. [9] 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. [10] 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. [11] has estimated per-pixel reflectance lobe widths by adding 2nd-order spherical harmonic illumination conditions to the basis of Ma et al. [10], 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 Ma et al. [10] actually *overdetermine* the surface albedo (one parameter per color) and surface normal (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 [12] 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 Debevec et al. [13]. 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.




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{v} towards a camera.

II. METHOD

We photograph the scene inside a geodesic sphere of colored LED lights, which can be programmed to produce gradient illumination as well as uniform white illumination, similar to Ma et al. [10]. We compute the mean spherical angle of reflected light and the total amount of reflected light as in Ma et al. [10]. However, we make additional use of the information contained within the denominator in the normalization step of computing the mean spherical angle. This allows us to solve for an additional parameter in the reflectance function, which in our case is the exponent n of a cosine lobe reflectance function.

We explore two different cosine lobe reflectance functions, which integrate easily over the uniform and gradient illumination conditions, so that an analytic solution for the fit 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, k and n are constants.

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.

A. Data Acquisition

In the data acquisition stage, we make four observations o_0, 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 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 ratio image of the colored gradient image to the full-on image, motivated by Woodham [12]. This process is illustrated in Fig. 3. Also, 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.




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 (B) illumination conditions. Middle: The red, green, and blue components of the ratio of (B) to (A). Bottom: The middle row multiplied by (A).

B. Fitting Cosine Lobe Reflectance Functions

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

axis of the lobe, which 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:

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

Substituting the four observations o_0, 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:

$$\begin{aligned} o_0 &= \frac{2\pi k L}{n+1}, \\ o_x &= \frac{\pi k L \hat{\alpha}_x}{n+2} + \frac{\pi k L}{n+1}, \\ o_y &= \frac{\pi k L \hat{\alpha}_y}{n+2} + \frac{\pi k L}{n+1}, \\ o_z &= \frac{\pi k L \hat{\alpha}_z}{n+2} + \frac{\pi k L}{n+1}, \\ \|\hat{\alpha}\| &= 1. \end{aligned}$$

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

$$\begin{aligned} \hat{\alpha} &= \frac{\alpha}{\|\alpha\|}, \quad (2a) \\ n &= \frac{2\|\alpha\| - o_0}{o_0 - \|\alpha\|}, \quad (2b) \\ k &= \frac{o_0(n+1)}{2\pi L}, \quad (2c) \end{aligned}$$

where $\alpha = \langle 2o_x - o_0, 2o_y - o_0, 2o_z - o_0 \rangle$.

Note that the cosine lobe axis is obtained from the gradient illumination condition observations in a manner similar to the method in Ma et al. [10] 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.

It can be shown by a similar argument that a spherical cosine lobe of the form $k(\frac{1}{2}\hat{\alpha} \cdot \Theta + \frac{1}{2})^n$ has the solution:

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

$$n = \frac{2\|\alpha\|}{o_0 - \|\alpha\|}, \quad (3b)$$

$$k = \frac{o_0(n+1)}{4\pi L}. \quad (3c)$$



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 one eighth.



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.

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 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.

C. Estimating Geometry and Shadows

The cosine lobe reflectance function $f(\Theta)$ estimated by our method has only low-frequency features. Thus, while the low-frequency illumination condition input photographs are reconstructed exactly, hard cast shadows cannot be synthesized for novel illumination conditions having high-frequency features. In order to enable the computation of hard cast shadows in the relit images, we estimate the scene geometry by integrating photometric surface normals. For each pixel, the surface normal $\hat{\mathbf{n}}$ is estimated heuristically from the fitted cosine lobe reflectance function, interpolating between the axis of the lobe for diffuse materials, and the half-angle between the view vector $\hat{\mathbf{v}}$ and the axis of the lobe for specular materials:

$$\hat{\mathbf{n}} \approx \frac{\hat{\alpha} + \hat{\mathbf{v}} \max(0, \min(1, n-1))}{\|\hat{\alpha} + \hat{\mathbf{v}} \max(0, \min(1, n-1))\|}. \quad (4)$$

Surfaces in the distant background that fall outside our illumination sphere receive very little light, and exhibit noise in the estimated normals that would severely warp the integrated geometry. To reduce this problem, we interpolate the estimated normal towards the z axis when the quantity $\|\alpha\|$ falls below a small epsilon. Finally, we blend the normals from the red, green, and blue channels to obtain our surface normal estimate. Once the normals are estimated, normal integration proceeds using a method similar to Horn et al. [14]. See Fig. 6 and 7 for examples of reconstructed geometry.

Once the scene geometry has been estimated, we compute an adjusted reflectance function $f'(\Theta)$ that exhibits hard shadows cast by the geometry. First, consider a reflectance function $f^*(\Theta)$ adjusted simply by occluding any light directions that hit scene geometry:

$$f^*(\Theta) = f(\Theta) \text{visible}(\Theta), \quad (5)$$

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 generally be altered by the visibility masking. To correct this, we introduce a scaling factor s to restore the overall reflectance magnitude, and introduce a low-frequency ‘‘ambient’’ term $p \cdot \Theta$ to restore




Figure 6. Reconstructed geometry from photometric normals. Top: Two views of the cat subject. Bottom: Two views of the duck subject.




Figure 7. Reconstructed geometry for the cat subject, with normals derived from two photos instead of four.

the original low-frequency component of the reflectance function:

$$f'(\Theta) = sf^*(\Theta) + p \cdot \Theta. \quad (6)$$

Let o_0^* , o_x^* , o_y^* , o_z^* represent the synthesized observations using $f^*(\Theta)$, computed by numerical integration:

$$\begin{aligned} o_0^* &= \int_{S^2} f^*(\Theta) L d\omega_\Theta, \\ o_x^* &= \int_{S^2} f^*(\Theta) \left(\frac{1}{2}\Theta_x + \frac{1}{2}\right) L d\omega_\Theta, \\ o_y^* &= \int_{S^2} f^*(\Theta) \left(\frac{1}{2}\Theta_y + \frac{1}{2}\right) L d\omega_\Theta, \\ o_z^* &= \int_{S^2} f^*(\Theta) \left(\frac{1}{2}\Theta_z + \frac{1}{2}\right) L d\omega_\Theta. \end{aligned}$$

Then we wish to correct the observations:

$$\begin{aligned} o_0 &= so_0^* + \int_{S^2} (p \cdot \Theta) L d\omega_\Theta, \\ o_x &= so_x^* + \int_{S^2} (p \cdot \Theta) \left(\frac{1}{2}\Theta_x + \frac{1}{2}\right) L d\omega_\Theta, \\ o_y &= so_y^* + \int_{S^2} (p \cdot \Theta) \left(\frac{1}{2}\Theta_y + \frac{1}{2}\right) L d\omega_\Theta, \\ o_z &= so_z^* + \int_{S^2} (p \cdot \Theta) \left(\frac{1}{2}\Theta_z + \frac{1}{2}\right) L d\omega_\Theta, \end{aligned}$$

which admit one closed-form solution:

$$s = \frac{o_0}{o_0^*}, \quad (7a)$$

$$p = \frac{3}{4\pi L} \langle o_x - so_x^*, o_y - so_y^*, o_z - so_z^* \rangle. \quad (7b)$$

The plausibility of the low-frequency term $p \cdot \Theta$ 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. 8. Indirect bounce light is often low-frequency in nature, so it seems appropriate to model it by a low-frequency term.




Figure 8. 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 s and p , 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. 9 shows such a mask, and compares the results obtained with and without using the mask.




Figure 9. 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.

III. CONCLUSIONS AND FUTURE WORK

Refer to Fig. 10 and 11 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 shows a noticeable reduction in quality compared to the four-photograph hemispherical cosine lobe method, most notably in areas of high color saturation such as the orange eyes, but could be suitable for applications with limited capture budget such as video relighting. The four-photograph hemispherical cosine lobe method appears to work well for most materials on the cat. Notably, the cosine lobe exponents visualized in Fig. 4 largely agree with the materials of the actual subject. For example, the golden body of the cat has exponents ranging from 1.2 up to 6.0 at grazing angles, indicating a specular material with dependency on angle, in this case a gold paint. The red collar, ears, and nose have exponents of 1.1 in the red channel, and higher in the green and blue channels, indicating 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 exponents of 0.7 in the green channel, and somewhat higher in the red and blue channels, in this case a very soft diffuse green paint with a faint gloss. Exceptions include highly specular materials, which are noisy due to the hardware limitation of using LED lights to approximate continuous illumination, and softening of hard cast shadows. It is clear that the estimated geometry is not an accurate reconstruction of the scene, but in a relighting scheme that does not attempt to alter the

viewpoint, the hard cast shadows and self shadowing from the reconstructed geometry 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 conditions 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. 12, 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.

Ground truth comparisons reveal several limitations with our method, but even with these limitations the method produces visually plausible results over a wide range of materials and illumination using very few input observations. Furthermore, the results are more consistent with ground truth than those obtained from using the same input photographs as a linear basis.



Figure 12. The visibility masking assumption does not hold for highly scattering materials such as the duck subject, resulting in an inappropriately hard shadow terminator.

ACKNOWLEDGEMENTS

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

REFERENCES

- [1] P. Haeberli, “Synthetic lighting for photography,” January 1992, available at <http://www.graficaobscura.com/synth/>.

- [2] J. S. Nimeroff, E. Simoncelli, and J. Dorsey, "Efficient re-rendering of naturally illuminated environments," in *Fifth Eurographics Workshop on Rendering*, Jun. 1994, pp. 359–373.
- [3] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar, "Acquiring the reflectance field of a human face," in *Proceedings of ACM SIGGRAPH 2000*, ser. Computer Graphics Proceedings, Annual Conference Series, Jul. 2000, pp. 145–156.
- [4] P.-P. Sloan, J. Kautz, and J. Snyder, "Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments," *ACM Transactions on Graphics*, vol. 21, no. 3, pp. 527–536, July 2002.
- [5] R. Ramamoorthi and P. Hanrahan, "Frequency space environment map rendering," *ACM Transactions on Graphics*, vol. 21, no. 3, pp. 517–526, July 2002.
- [6] R. Wang, J. Tran, and D. Luebke, "All-frequency interactive relighting of translucent objects with single and multiple scattering," *ACM Transactions on Graphics*, vol. 24, no. 3, pp. 1202–1207, Aug. 2005.
- [7] R. Ramamoorthi, "Modeling Illumination Variation with Spherical Harmonics," in *Face Processing: Advanced Modeling Methods*, 2006, pp. 385–424.
- [8] A. Wenger, A. Gardner, C. Tchou, J. Unger, T. Hawkins, and P. Debevec, "Performance relighting and reflectance transformation with time-multiplexed illumination," *ACM Transactions on Graphics*, vol. 24, no. 3, pp. 756–764, Aug. 2005.
- [9] T. Malzbender, D. Gelb, and H. Wolters, "Polynomial texture maps," in *Proceedings of SIGGRAPH 2001*, pp. 519–528.
- [10] W.-C. Ma, T. Hawkins, P. Peers, C.-F. Chabert, M. Weiss, and P. Debevec, "Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination," in *Rendering Techniques 2007: 18th Eurographics Symposium on Rendering*, Jun. 2007, pp. 183–194.
- [11] A. Ghosh, T. Chen, P. Peers, C. A. Wilson, and P. Debevec, "Estimating specular roughness and anisotropy from second order spherical gradient illumination," in *Rendering Techniques 2009: 20th Eurographics Symposium on Rendering*, Jun. 2009.
- [12] R. J. Woodham, "Photometric method for determining surface orientation from multiple images," *Optical Engineering*, vol. 19, no. 1, pp. 139–144, 1980.
- [13] P. Debevec, A. Wenger, C. Tchou, A. Gardner, J. Waese, and T. Hawkins, "A lighting reproduction approach to live-action compositing," *ACM Transactions on Graphics*, vol. 21, no. 3, pp. 547–556, Jul. 2002.
- [14] B. K. P. Horn and M. J. Brooks, "The variational approach to shape from shading," vol. 33, 1986, pp. 174–208.



Figure 10. Relit images of cat subject under novel point light illumination conditions. A) Ground truth. B) Linear basis using four photographs. C) Hemispherical cosine lobes using two photographs. D) Hemispherical cosine lobes using four photographs. E) Hemispherical cosine lobes with hard cast shadows, using four photographs and user-supplied visibility mask.



Figure 11. Relit images of duck subject under novel point light illumination conditions. A) Ground truth. B) Linear basis using four photographs. C) Spherical cosine lobes using two photographs. D) Spherical cosine lobes using four photographs.