Determining Light Direction in Spheres
using Average Gradient

Richard Dosselmann and Xue Dong Yang
Technical Report TR-CS 2009-1
February 2009

Department of Computer Science
University of Regina
Regina, Saskatchewan, Canada S4S 0A2

ISBN 978-0-7731-0665-9 (print)
ISBN 978-0-7731-0666-6 (on-line)
ISSN 0828-3494

Determining Light Direction in Spheres using Average Gradient

Richard Dosselmann and Xue Dong Yang
Department of Computer Science
University of Regina
Regina, Saskatchewan, Canada S4S 0A2
{dosselmr, yang}@cs.uregina.ca

Abstract

Using only the individual gradient vectors of a shaded image of an illuminated sphere, the position of the light source in a scene is computed. A two-part process, the proposed method determines both the two-dimensional planar direction of the light source, as well as its angular tilt above the image plane. Uncomplicated and therefore easily implemented, the method boasts an exceptionally fast run-time. It is currently designed for use in scenes involving a single point light source and a single shaded sphere, though it readily extends to ellipsoids, and, at times, textured spheres. In the case of arbitrary surfaces, only the two-dimensional planar direction of the light source may be estimated at this time. A formal proof of the proposed method, up to the two-dimensional planar direction of the light source, is also presented in this work.

Keywords: Light direction, shape-from-shading, average gradient, tilt angle, sphere, ellipsoid

1 Introduction

Light direction estimation is frequently considered as a first step in the solution of a general *shape-from-shading* [1, 2, 3] problem. Simply stated, the objective is to determine the three-dimensional position of a light source in a scene, given a two-dimensional image of that scene. Because three-dimensional information is normally lost in the production of an image, one is not generally able to accurately recover the position of a light source from a single image. This paper proposes a straightforward and fast solution to this problem in the case of a shaded sphere under a single light source. Specifically, it has been discovered that the average gradient vector of the shaded image, along with its magnitude, is sufficient to accurately determine the position of the light source, up to a scalar factor of its distance from the scene. This finding is proved, up to the two-dimensional planar direction of the light source, in Appendix B. While the technique is largely aimed at smoothly-shaded *Lambertian* [4] spheres, it partially extends to textured surfaces, ellipsoids and even general scenes.

A number of motivating factors inspire research in this area, not the least of which is that of a stepping stone to a solution to the more general problem of shape-from-shading and scene reconstruction. Additional applications include interactive photo montage [5], facial recognition under changing lighting conditions [6], the removal of specular highlights [7] and three-dimensional robot vision [8].

Light detection is classic problem in image processing. Among the earliest of all works in this area is that of Brooks and Horn [1]. In this paper, the two present analytic solutions to both the light direction and shape determination problems, without explicit use of the image *reflectance map* [9]. A number of analytic approaches [9, 10] involving the reflectance map, and in particular image contours, have since appeared. Using an alternate approach, Debevec [11] estimates local lighting information by positioning a mirrored sphere into a given scene. However, this method requires that the scene geometry be known. This highly restrictive assumption is also made by Upright et al. in [12], in which the group uses wavelets to recover lighting information. Next, assuming that the geometry of the object of interest is known, Hara et al. [13] are able to estimate the position of a light source using a single image of the given object. By later varying the position of the light, the team is able to generate any number of synthetic images of the object. In [14], Matsushita et al. approximate the illumination of a scene by way of a minimization process. A further listing and overview of techniques, both classic and modern, is presented in [15]. Although the aforementioned techniques may be applied to spheres, they are largely designed for general surfaces. One of the first works to specifically address spheres is that of Lee and Rosenfeld [16], published in 1985. Later,

using neural networks, Dror et. al [17] examine the surface reflectance properties of a variety of spheres. Zhou and Kambhamettu [18] establish the direction and intensity of several light sources using two stereo images of a given sphere. An iterative routine using least squares and various critical points is presented in [19]. In yet another attempt, Yang and Yuille [20] decide among a number of probable light positions, based on the shading conditions present in the given photograph.

The new method is introduced and described in Section 2. Then, in Section 3, it is extended to other surfaces, followed by some testing carried out in Section 4, before concluding remarks are made in Section 5.

2 Light Direction Detection

The three-dimensional direction of a light source, represented by the normalized vector $\mathbf{L} = \langle L_x, L_y, L_z \rangle$, is obtained via a two-stage process. In the first step, the average gradient vector of the shaded image is computed. Then, using the magnitude of this average vector, the tilt angle of the light above the image plane is calculated. Together, these two components accurately identify the three-dimensional direction of the light source from the scene. This entire procedure requires only a single shaded image of an illuminated sphere. Unfortunately, the approach does not allow one to determine the precise distance of the light from the scene. As a result, $\mathbf{L}$, as indicated above, is given as a normalized vector, even though the actual light source may not be at an exact unit distance from the scene. Psuedocode for the complete algorithm is found in Appendix A.

2.1 Average Gradient

This research finds that the average of the two-dimensional gradient vectors of a given shaded image of a sphere is equal to the vector that points in the direction of the light source illuminating the scene. This remarkable fact is a result of the elegant symmetrical properties of the gradient vectors in a shaded image of a sphere. Consider, for instance, the example of Fig. 1. Here, the light is located to the upper right of the scene, as is evident from the shaded sphere of Fig. 1(a). Not surprisingly, a large number of the gradient vectors in the corresponding gradient map of Fig. 1(b), specifically those in the central region, point in this direction. Those near the top left of the map, however, generally point in a horizontal direction, while those in the lower right section are directed upwards. When averaged, these two sets of horizontal and vertical vectors also produce a vector oriented to the upper right and in the direction of the light source. Further observe that these two sets are symmetric with respect to the direction of the light. It appears that for each vector on one side of the line passing through the light, there is a complementary vector on the other side. The average of each pair, in general, is a vector in the direction of the light. When these vectors are later averaged with those in the central portion of the gradient map, the ensuing global average is a two-dimensional vector aimed squarely at the light. Though quite simple, this process is surprisingly adept at calculating the direction of the light source. This fact is formally established in Appendix B.

Mathematical descriptions of the local and global gradient vectors are given below. Note that in order to later determine the vertical tilt angle of the light source, each of the local gradient vectors must be individually normalized, as further discussed in Section 2.2. Assume that the gray-level shading value of each image pixel is represented by the function $f(x, y)$, where x indicates the column of the pixel and y identifies its row. The origin, occurring in the upper-left corner of the image, is represented by the point $(0, 0)$. From here, x increases in a positive fashion from left to right, while y increases in a likewise manner from top to bottom. The local *gradient* [21] vector at any pixel $f(x, y)$ is then computed as

$$\nabla f(x, y) = \langle f_x(x, y), f_y(x, y) \rangle = \langle f(x, y) - f(x - 1, y), f(x, y) - f(x, y - 1) \rangle \,, \tag{1}$$

where $f_x(x, y)$ and $f_y(x, y)$ denote the horizontal and vertical components, respectively, of the gradient. It is important to recognize that these gradients are computed across the shaded image of the sphere rather than the actual surface of the sphere itself. In addition, because of the discrete boundaries of a digital

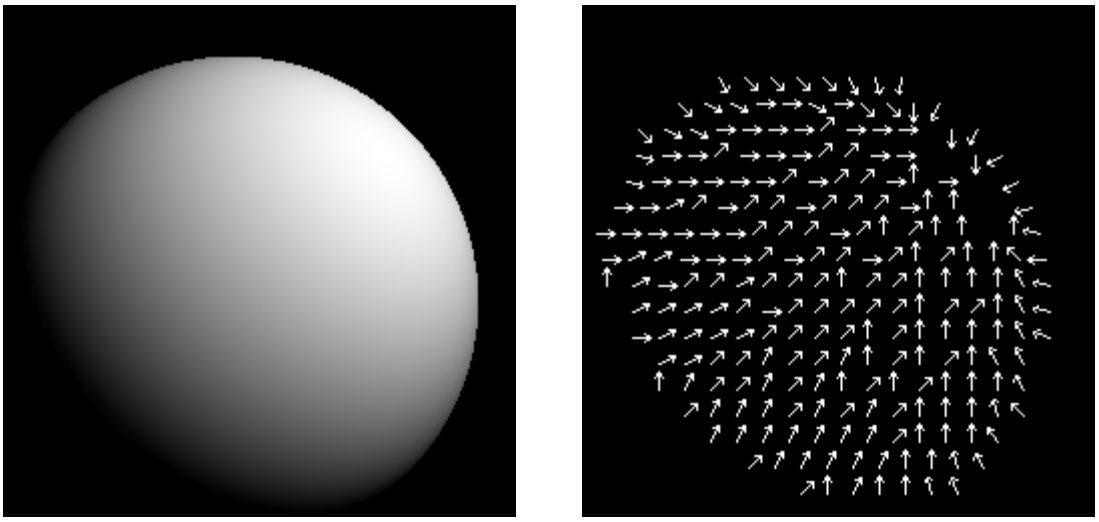

(a) Shaded sphere with light at $\mathbf{L} = \langle 2, 2, 3 \rangle$ (tilt angle of 47°)

(b) Gradient map of shaded image of sphere

Figure 1: Shaded sphere and corresponding gradient map.

image, it is not possible to compute a gradient vector for those pixels lying along the image border. The normalized form of each local gradient vector is given as

$$\nabla f^*(x,y) = \left\langle f_x^*(x,y), f_y^*(x,y) \right\rangle = \left\langle \frac{f_x(x,y)}{\|\nabla f(x,y)\|}, \frac{f_y(x,y)}{\|\nabla f(x,y)\|} \right\rangle. \tag{2}$$

The individual horizontal and vertical components are averaged over all local gradient vectors to produce a single global average gradient, given as

$$\nabla \bar{f}(x,y) = \left\langle \bar{f}_x(x,y), \bar{f}_y(x,y) \right\rangle = \left\langle \frac{1}{n} \sum_{i=1}^{n} f_x^*(x,y), \frac{1}{n} \sum_{i=1}^{n} f_y^*(x,y) \right\rangle, \tag{3}$$

where n represents the number of non-zero gradients in the picture. At present, a non-zero gradient is one in which the absolute value of at least one of its horizontal or vertical elements exceeds a threshold of 0.01. Furthermore, as a primitive means of excluding object boundaries from the computation of the average gradient, local gradient vectors in which either of the horizontal or vertical components exceeds 32 are ignored. Note that this entire procedure requires only a single pass over the input image. In most instances, the algorithm runs in only a few milliseconds.

2.2 Tilt Angle

Closer inspection of the gradient map of Fig. 1(b) reveals that the vectors in the extreme upper right ribbon-shaped region of the sphere are in fact oriented in the opposite direction of the light source. Averaging these vectors with the others, $\nabla \bar{f}(x,y)$ is shortened accordingly. The degree of shortening is proportionate to the number of these "opposing" vectors. And the number of these vectors is related to the vertical tilt angle of the light, depicted in Fig. 2. As a rule, the greater the number of these "opposing" vectors, the higher the tilt angle of the light source. When there are no such "opposing" vectors, $\|\nabla \bar{f}(x,y)\| = 1.0$. This corresponds to a tilt angle of $\theta = 0°$, meaning that the light lies in the image plane. In the exact opposite situation, in which the combined length of the opposing vectors is equal to that of those pointed at the light, $\|\nabla \bar{f}(x,y)\|$ is reduced to 0.0. This time, the light is positioned directly overhead, perpendicular to the plane, with $\theta = 90°$. In general, the tilt angle is given as

$$\theta = \arccos\left(\|\nabla \bar{f}(x,y)\|\right), \tag{4}$$

where $0° \leq \theta \leq 90°$ and $0 \leq \|\nabla \bar{f}(x,y)\| \leq 1$. Angles are strictly positive, meaning that light sources behind the image plane are not permitted. As well, $\|\nabla \bar{f}(x,y)\|$ never exceeds 1.0 since each of the local vectors is individually normalized.

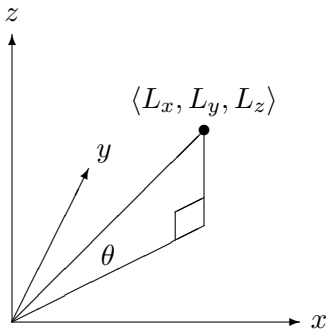

Figure 2: Tilt angle.

It was quickly determined that the individual gradient vectors needed to be normalized in order to obtain an accurate estimate of the tilt angle. Though not yet fully understood, this idea appears to be related to the manner in which local vectors not directly oriented at the light source "cancel out" one another. This phenomenon is best illustrated by way of an example. Take again the situation of Fig. 1(b). Specifically, consider the average of one horizontal vector, given as $\langle 1.0, 0.0 \rangle$, and one vertical gradient vector, namely $\langle 0.0, 1.0 \rangle$. Having been normalized, each of these vectors has a length of 1.0. By (3), the average of the two is found to be

$$\nabla \bar{f}(x,y) = \left\langle \frac{1}{2}(1.0 + 0.0), \frac{1}{2}(0.0 + 1.0) \right\rangle \tag{5}$$

$$= \left\langle \frac{1}{2}, \frac{1}{2} \right\rangle. \tag{6}$$

The length of $\nabla \bar{f}(x,y)$ is equal to

$$\|\nabla \bar{f}(x,y)\| = \sqrt{\left(\frac{1}{2}\right)^2 + \left(\frac{1}{2}\right)^2} = \frac{1}{\sqrt{2}} < 1. \tag{7}$$

As is obvious from (7), the length of the average vector is less than 1.0. Conversely, those vectors aimed directly at the light source have a length very close to 1.0. Suppose that these vectors are averaged with those of the sort given in the example above. In such a case, the length of the resulting average vector is less than 1.0. This decrease in length increases the estimate of the tilt angle. Moreover, as the light rises higher above the scene, there are increasingly more horizontal and vertical vector pairs, further reducing the length of the global average vector and increasing the value of the tilt angle. It appears that the amount of reduction is directly related to the actual value of the tilt angle. Additional research is needed to precisely explain this phenomenon. In another sense, by normalizing the vectors, the global average essentially becomes a count of the number of vectors pointed at that light versus those aimed in the opposite direction. Each of these ideas warrants further study. One point that is clear, however, is that without such normalization, the resulting tilt angle is often noticeably inaccurate.

A slight error in the measure of the tilt angle will periodically occur in situations involving small tilt angles. In such cases, the contour lines produced on a sphere by the lighting are oriented in directions perpendicular to the boundaries of the sphere. As a result, local gradient vectors $\nabla f(x,y)$ are often oriented away from the light source. Consequently, $\nabla \bar{f}(x,y)$ is shortened and therefore causes the tilt angle to appear higher than in actual fact. For larger angles, contour lines are oriented in the direction

of the light source and produce a longer average gradient vector and thus more accurate estimates of θ. Note that not all small angle measures result in errors. Further studies of the exact conditions that must be met in order to observe this error are planned. As an example, see Fig. 3(a) of Appendix C.

3 Additional Surfaces

In the following sections, the method of Section 2 is extended to situations involving textured spheres, ellipsoids and even general surfaces and photographic scenes. In the case of the textured spheres and ellipsoids, it is still possible to estimate the tilt angle of the light. Among the remaining surfaces, though, this is generally not true.

3.1 Textured Spheres

Textures pose a challenge to the algorithm in its existing form. It is not possible at this time to distinguish between image contours caused by lighting versus those resulting from the shading and patterns of the texture. An average of all image contours, $\nabla \bar{f}(x, y)$ is sometimes incorrect. Perhaps the simplest means of addressing this issue is to blur the input image in order to diminish the effects of the textures. This is the approach taken in the examples of Fig. 7 of Appendix C. Another solution involves lowering the "segmentation" threshold of Section 2.1, currently set at 32. Such a change may result in the algorithm missing various lighting contours, though the impact is likely to be marginal.

3.2 Ellipsoids

A scaled variation of the sphere, the *ellipsoid* [21] is mathematically defined as

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} + \frac{z^2}{c^2} = 1, \tag{8}$$

where $a > 0$, $b > 0$ and $c > 0$ denote the radii along each of the primary axes. Assuming a parallel viewing projection along the z-axis, the ellipsoid appears as a two-dimensional ellipse of the form

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1. \tag{9}$$

Suppose, without loss of generality, that the horizontal radius of the ellipse is the longer of the two radii. The length of this longer radius is taken to have unit length, that is $a = 1$. The ellipse of (9) is therefore simplified to

$$x^2 + \frac{y^2}{b^2} = 1. \tag{10}$$

The radius of the shorter, vertical side is equal to the length of that radius, divided by the actual length of the horizontal radius. It is expressed as

$$0 \leq \frac{b}{a} = b \leq 1. \tag{11}$$

Observe that if one were to scale the second term of the expression of (10) by b^2, the ellipse would be reduced to a circle. Hence, the computation of each local gradient vector is modified from that of (1) to the following scaled form, namely

$$\nabla f(x, y) = \langle f_x(x, y), f_y(x, y) \rangle = \left\langle f(x, y) - f(x - 1, y), b^2 \left(f(x, y) - f(x, y - 1) \right) \right\rangle. \tag{12}$$

A more formal investigation of this idea is expected at a later date. Examples involving ellipsoids are given in Figs. 8 and 9 of Appendix C.

3.3 Arbitrary Surfaces

An arbitrary surface may represent any surface or combination of surfaces. Lifelike scenes composed of textures, shadows, detailed objects and complex reflections also fall into this category. Though $\nabla \bar{f}(x, y)$ is usually a strong indication of the direction of the light source, it is occasionally less accurate than in the previous situations. Furthermore, there is no general means of estimating the tilt angle in these sorts of scenarios. The examples of Fig. 10 of Appendix C show the performance of the algorithm in these instances.

4 Experiments

A large number of tests involving synthetic and actual images of objects have been carried out. Artificial images were rendered using OpenGL®. All examples are found in Appendix C. Examples of Lambertian spheres are given in Figs. 3, 4, 5 and 6. In each, the measure of $\nabla \bar{f}(x, y)$ is very reasonable. Moreover, the estimated tilt angles are extremely close to their corresponding actual values. Smaller angles, most notably that of Fig. 3(a), are less accurate. Please refer to Section 2.2 for a discussion of this matter. An assortment of textured spheres is seen in Fig. 7. Blurring was achieved using a 2.0−radius Gaussian blur. This simple heuristic appears to be rather effective. In each of the four images shown, $\nabla \bar{f}(x, y)$ clearly points in the direction of the light, though the precise direction of the source is not available. Examples of illuminated ellipsoids are given in Figs. 8 and 9. In each of the four ellipsoids of Fig. 8, the vertical height of the object is 80% that of its width. Therefore, a scaling value of $b^2 = 0.8^2 = 0.64$ is employed, as described in Section 3.2. Whereas the calculated values of $\nabla \bar{f}(x, y)$ are relatively accurate, the estimates of the tilt angle are somewhat less so, particularly when compared with those of the spheres. In the second example, namely that of Fig. 9, a scaling factor of $b^2 = 0.6^2 = 0.36$ is required. Natural scenes are given in Fig. 10. In this example, like that of Fig. 7, the actual position of the light source is not available. The computed values of $\nabla \bar{f}(x, y)$ nevertheless point towards the illuminating sources.

5 Conclusions

The method presented in this paper offers a swift, yet reliable, answer to the light detection problem when the object of interest is a sphere. Most appealing of all, the algorithm involves nothing more than the gradient vectors of a single shaded image. From here, a straightforward calculation allows one to approximate the vertical tilt angle of the light source.

A number of enhancements to extend the ideas of this research to shapes and surfaces beyond spheres and ellipsoids are actively being explored. Such ideas include, for instance, assuming arbitrary surfaces to be locally spherical, blurring images to reduce the effects of object boundaries or using segmentation to separate objects from one another. Before doing this, however, it is important that the existing approach be broadened to deal with images that include ambient background light. This matter is perhaps easily handled when it is known in advance that the scene contains a sphere. The strategy then is to consider those regions in which the local gradient vector $\nabla f(x, y)$ is near zero, while being careful to ignore background areas outside of the sphere. These low-detail regions with few contours are likely to free of the effects of diffuse lighting. Consequently, any shading information must be the result of ambient lighting. The pixels in such regions naturally offer a reasonable estimate of the amount of ambient light in the scene. Once this estimate of ambient light is subtracted from the image, the method of Section 2 may be called upon to determine the position of the light.

References

[1] Michael J. Brooks and Berthold K. P. Horn. Shape and source from shading. *Massachusetts Institute of Technology Artificial Intelligence Laboratory*, January 1985.

[2] Bertrand Kerautret, Xavier Granier, and Achille Braquelaire. Intuitive shape modeling by shading design. *Lecture Notes in Computer Science*, 3638:163–174, 2005.

[3] Pawan Sinha and Edward Adelson. Recovering reflectance and illumination in a world of painted polyhedra. *Fourth Int. Conf. Computer Vision*, pages 156–163, May 1993.

[4] James D. Foley, Andries van Dam, Steven K. Feiner, John F. Hughes, and Richard L. Phillips. *Introduction to Computer Graphics*. Addison-Wesley Publishing Company, 1994.

[5] Aseem Agarwala, Mira Dontcheva, Maneesh Agrawala, Steven Drucker, Alex Colburn, Brian Curless, David Salesin, and Michael Cohen. Interactive digital photomontage. *Proc. 2004 SIGGRAPH Conf.*, 23(3):294–302, August 2004.

[6] Kuang-Chih Lee and Baback Moghaddam. A practical face relighting method for directional lighting normalization. *Lecture Notes in Computer Science*, 3723:155–169, 2005.

[7] Satya P. Mallick, Todd Zickler, Peter N. Belhumeur, and David J. Kriegman. Specularity removal in images and videos: A pde approach. *Lecture Notes in Computer Science*, 3951:550–563, 2006.

[8] Miao Liao, Liang Wang, Ruigang Yang, and Minglun Gong. Light fall-off stereo. *IEEE Conf. Computer Vision and Pattern Recognition, 2007*, pages 1–8, June 2007.

[9] Achille Braquelaire and Bertrand Kerautret. Reconstruction of lambertian surfaces by discrete equal height contours and regions propagation. *Image and Vision Computing*, 23(2):177–189, February 2005.

[10] Ron Kimmel and Alfred M. Bruckstein. Global shape from shading. *Computer Vision and Image Understanding*, 62(3):360–369, November 1995.

[11] Paul Debevec. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. *Proc. 25th Conf. Computer Graphics and Interactive Techniques*, pages 189–198, July 1998.

[12] Cameron Upright, Dana Cobzaş, and Martin Jägersand. Wavelet-based light reconstruction from a single image. *Proc. Fourth Canadian Conf. Computer and Robot Vision*, pages 305–312, 2007.

[13] Kenji Hara, Ko Nishino, and Katsushi Ikeuchi. Light source position and reflectance estimation from a single view without the distant illumination assumption. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 27(4):493–505, April 2005.

[14] Yasuyuki Matsushita, Stephen Lin, Sing Bing Kang, and Heung-Yeung Shum. Estimating intrinsic images from image sequences with biased illumination. *Lecture Notes in Computer Science*, 3022:274–286, 2004.

[15] Nathan Funk. A survey of light source detection methods. University of Alberta: Mini-Project for CMPUT 603, November 2003.

[16] Chia-Hoang Lee and Azriel Rosenfeld. Improved methods of estimating shape from shading using the light source coordinate system. *Artificial Intelligence*, 26(2):125–143, 1985.

[17] Ron O. Dror, Edward H. Adelson, and Alan S. Willsky. Estimating surface reflectance properties from images under unknown illumination. *Proc. SPIE 4299: Human Vision and Electronic Imaging IV*, 4299:231–242, January 2001.

[18] Wei Zhou and Chandra Kambhamettu. Estimation of illuminant direction and intensity of multiple light sources. *Lecture Notes in Computer Science*, 2353:206–220, 2002.

[19] Yufei Zhang and Yee-Hong Yang. Multiple illuminant direction detection with application to image synthesis. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 23(8):915–920, August 2001.

[20] Yibing Yang and Alan Yuille. Sources from shading. *Proc. CVPR '91 IEEE Computer Society Conf. Computer Vision and Pattern Recognition*, pages 534–539, June 1991.

[21] James Stewart. *Calculus*. Brooks/Cole Publishing Company, Pacific Grove, California, 1987.

[22] Robert H. Risch. The problem of integration in finite terms. *Trans. American Mathematical Society*, 139:167–189, May 1969.

[23] A. D. Fitt and G. T. Q. Hoare. The closed-form integration of arbitrary functions. *Mathematical Gazette*, pages 227–236, 1993.

Appendix A

Psuedocode of the two-stage light direction detection algorithm is given below.

$\bar{f}_x(x,y) \leftarrow 0$
$\bar{f}_y(x,y) \leftarrow 0$
$\nabla\bar{f}(x,y) \leftarrow 0$
$n \leftarrow 0$

for each pixel $f(x,y)$, $1 \leq x \leq width - 1$, $1 \leq y \leq height - 1$ **do**
    $f_x(x,y) \leftarrow f(x,y) - f(x-1,y)$
    $f_y(x,y) \leftarrow f(x,y) - f(x,y-1)$
    $\nabla f(x,y) \leftarrow \langle f_x(x,y), f_y(x,y) \rangle$

    if $(|f_x(x,y)| > 0.01$ **or** $|f_y(x,y)| > 0.01)$ **and** $(|f_x(x,y)| \leq 32$ **and** $|f_y(x,y)| \leq 32)$ **then**
        $f_x^*(x,y) \leftarrow \frac{f_x(x,y)}{\|\nabla f(x,y)\|}$
        $f_y^*(x,y) \leftarrow \frac{f_y(x,y)}{\|\nabla f(x,y)\|}$
        $\bar{f}_x(x,y) \leftarrow \bar{f}_x(x,y) + f_x^*(x,y)$
        $\bar{f}_y(x,y) \leftarrow \bar{f}_y(x,y) + f_y^*(x,y)$
        $\nabla\bar{f}(x,y) \leftarrow \nabla\bar{f}(x,y) + \langle \bar{f}_x(x,y), \bar{f}_y(x,y) \rangle$
        $n \leftarrow n + 1$
    end if
end for

if $n > 0$ **then**
    $\bar{f}_x(x,y) \leftarrow \frac{1}{n}\bar{f}_x(x,y)$
    $\bar{f}_y(x,y) \leftarrow \frac{1}{n}\bar{f}_y(x,y)$
    $\nabla\bar{f}(x,y) \leftarrow \frac{1}{n}\nabla\bar{f}(x,y)$
    $\theta \leftarrow \arccos\left(\|\nabla\bar{f}(x,y)\|\right)$
end if

Appendix B

This section includes a partial proof of the new light detection algorithm. At this time, a complete and formal proof of the ideas put forward in Section 2 is not feasible. Though such a proof remains a primary objective, efforts continue to be hindered by two significant issues. First, recall the relationship of Section 2.1, in which the global average vector is found to be proportionate to the two-dimensional direction of the light source. A rigorous evaluation of this global sum requires the mathematical integration of the individual local gradient vectors over the entire sphere. Although each such vector can readily be expressed in algebraic terms, the problems introduced by the complex quadratic expressions and radicals of the normalization term in the denominator present significant challenges from the perspective of integration. Successful resolution of this first dilemma rests on the ability to integrate the complicated expression of (25). With traditional techniques having yielded very little progress, attention has shifted to Risch's algorithm [22, 23] for generalized integration. Unfortunately, there is the matter of finding a complete and viable implementation of Risch's enormously complicated integration algorithm. Alternately, it may very

well be the case that no *elementary* [21] integral exists for this problem. The second issue of concern is the lack of a thorough understanding of the effects of normalizing local vectors, as described in Section 2.2. Whether the normalization of individual vectors is the appropriate action from a mathematical point of view or is merely a heuristic is not yet known. Until such time, it is premature to attempt to prove anything in this regard. The following theorem, at minimum, proves the existence of a relationship between the average of the local gradient vectors, without the effects of normalization, and the two-dimensional direction of the light source.

Theorem 1 *The average of the non-normalized local gradient vectors of a given image of a diffuse-shaded Lambertian sphere is proportionate to the two-dimensional spatial position of the light source illuminating the scene.*

Proof. Consider a sphere S of radius r

$$x^2 + y^2 + z^2 = r^2, \tag{13}$$

with surface normal

$$\mathbf{N} = \langle x, y, z \rangle . \tag{14}$$

Assume, without loss of generality, that S is centered in the scene. Further assume that the light illuminating the scene is represented by the normalized vector $\mathbf{L} = \langle L_x, L_y, L_z \rangle$. Using a diffuse, distance-independent lighting model [4], the local surface shading at any point (x, y, z) on S is computed. Then, using a parallel projection, S is projected onto a two-dimensional image plane $f(x, y)$. Here, $z = f(x, y)$ represents the shading or gray-level value of each pixel in the image, namely

$$\begin{aligned} f(x,y) &= (i_p k_d)(\mathbf{N} \cdot \mathbf{L}) && (15)\\ &= c_d (\mathbf{N} \cdot \mathbf{L}) && (16)\\ &= c_d (\langle x, y, z \rangle \cdot \langle L_x, L_y, L_z \rangle) && (17)\\ &= c_d \left(\left\langle x, y, \sqrt{r^2 - x^2 - y^2} \right\rangle \cdot \langle L_x, L_y, L_z \rangle \right) && (18)\\ &= c_d \left(x L_x + y L_y + \sqrt{r^2 - x^2 - y^2} L_z \right). && (19) \end{aligned}$$

In the continuous case, the gradient at any point on the image plane is computed as

$$\nabla f(x,y) = \left\langle \frac{\partial f(x,y)}{\partial x}, \frac{\partial f(x,y)}{\partial y} \right\rangle . \tag{20}$$

Accumulating the individual gradients over the span of the sphere, in particular from $-r$ to r, and subsequently dividing by n, one obtains

$$\begin{aligned} \nabla \bar{f}(x,y) &= \left\langle \frac{1}{n} \int_{-r}^{r} \frac{\partial f(x,y)}{\partial x}\, dx, \frac{1}{n} \int_{-r}^{r} \frac{\partial f(x,y)}{\partial y}\, dy \right\rangle && (21)\\ &= \left\langle \frac{1}{n} f(x,y) \Big|_{-r}^{r}, \frac{1}{n} f(x,y) \Big|_{-r}^{r} \right\rangle && (22)\\ &= \frac{2rc_d}{n} \langle L_x, L_y \rangle , && (23) \end{aligned}$$

where $\langle L_x, L_y \rangle$ denotes the two-dimensional planar direction of the light source. □

As a remark, had one actually attempted to prove the validity of the complete formulation involving normalization factors, they would have been confronted by the ominous problem of integrating the following,

namely

$$\nabla \bar{f}(x,y) = \left\langle \frac{1}{n}\int_{-r}^{r} \frac{\frac{\partial f(x,y)}{\partial x}}{\|\nabla f(x,y)\|}\, dx, \frac{1}{n}\int_{-r}^{r} \frac{\frac{\partial f(x,y)}{\partial y}}{\|\nabla f(x,y)\|}\, dy \right\rangle \tag{24}$$

$$= \left\langle \frac{1}{n}\int_{-r}^{r} L_x - \frac{\frac{x}{\sqrt{r^2-x^2-y^2}}}{\sqrt{L_x^2+L_y^2-\frac{2xL_x+2yL_y}{\sqrt{r^2-x^2-y^2}}+\frac{x^2+y^2}{r^2-x^2-y^2}}}\, dx, \right.$$

$$\left. \frac{1}{n}\int_{-r}^{r} L_y - \frac{\frac{y}{\sqrt{r^2-x^2-y^2}}}{\sqrt{L_x^2+L_y^2-\frac{2xL_x+2yL_y}{\sqrt{r^2-x^2-y^2}}+\frac{x^2+y^2}{r^2-x^2-y^2}}}\, dy \right\rangle. \tag{25}$$

Appendix C

A number of test images of Lambertian and textured spheres, along with several ellipsoids and various natural scenes, are presented in the following pages. The particular choice of lightning varies across the different examples. In some instances, the viewing direction of the scene is perpendicular to the camera, while in others the camera is positioned at an oblique angle to the object. Lastly, θ^* represents the actual tilt angle, while θ, as defined in Section 2.2, denotes its estimate using the new algorithm.

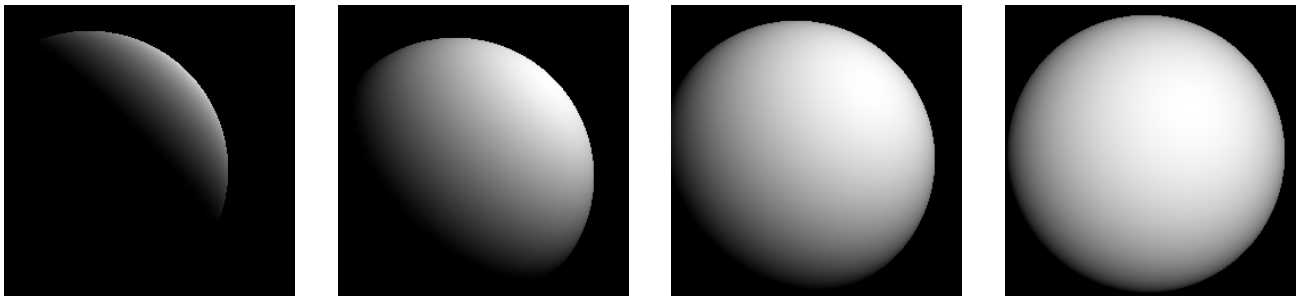

(a) $\nabla \bar{f}(x,y) = \langle 0.6551, -0.6798 \rangle, \theta^* = 2^\circ,\ \theta = 19.2616^\circ$

(b) $\nabla \bar{f}(x,y) = \langle 0.5722, -0.5724 \rangle, \theta^* = 35^\circ,\ \theta = 35.9659^\circ$

(c) $\nabla \bar{f}(x,y) = \langle 0.4036, -0.4082 \rangle, \theta^* = 55^\circ,\ \theta = 54.9707^\circ$

(d) $\nabla \bar{f}(x,y) = \langle 0.2501, -0.2471 \rangle, \theta^* = 71^\circ,\ \theta = 69.4164^\circ$

Figure 3: Lambertian spheres I.

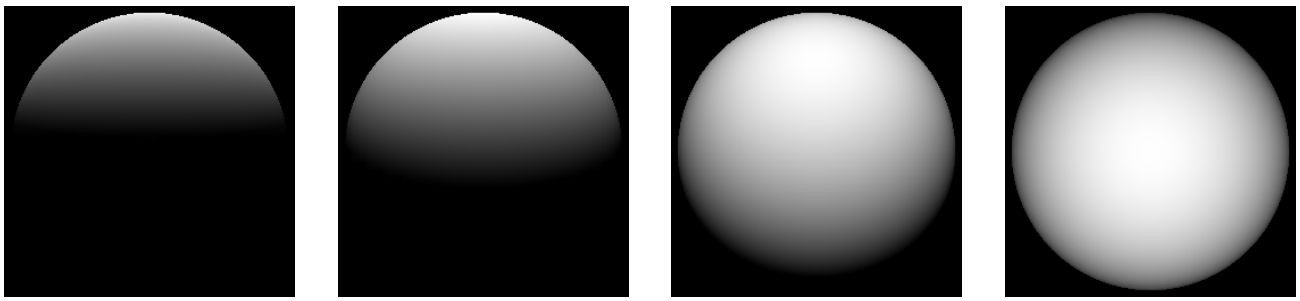

(a) $\nabla \bar{f}(x,y) = \langle 0.0061, -0.9860 \rangle, \theta^* = 7^\circ,\ \theta = 9.5891^\circ$

(b) $\nabla \bar{f}(x,y) = \langle 0.0247, -0.9236 \rangle, \theta^* = 22^\circ,\ \theta = 22.4859^\circ$

(c) $\nabla \bar{f}(x,y) = \langle 0.0106, -0.5009 \rangle, \theta^* = 60^\circ,\ \theta = 59.9352^\circ$

(d) $\nabla \bar{f}(x,y) = \langle 0.0020, 0.0049 \rangle, \theta^* = 90^\circ,\ \theta = 89.6962^\circ$

Figure 4: Lambertian spheres II.

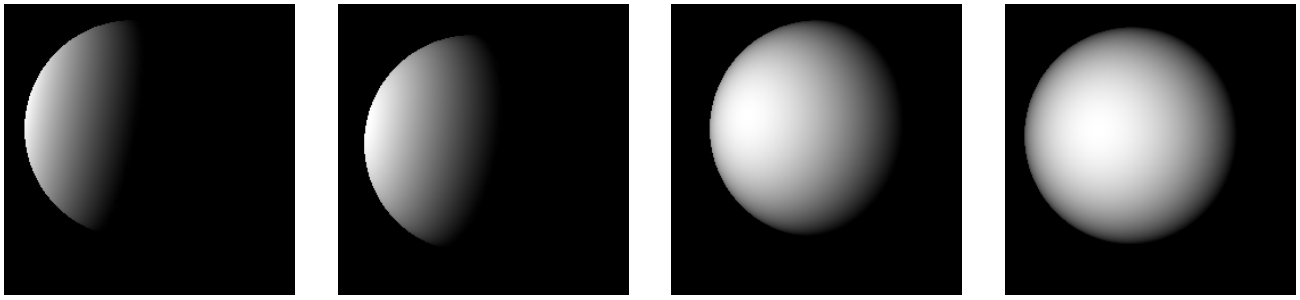

(a) $\nabla\bar{f}(x,y) = \langle -0.9524, -0.1484\rangle, \theta^* = 14^\circ,\ \theta = 15.4437^\circ$

(b) $\nabla\bar{f}(x,y) = \langle -0.8808, -0.1438\rangle, \theta^* = 25^\circ,\ \theta = 26.8173^\circ$

(c) $\nabla\bar{f}(x,y) = \langle -0.4807, -0.0791\rangle, \theta^* = 58^\circ,\ \theta = 60.8426^\circ$

(d) $\nabla\bar{f}(x,y) = \langle -0.2382, -0.0383\rangle, \theta^* = 75^\circ,\ \theta = 76.0379^\circ$

Figure 5: Lambertian spheres III.

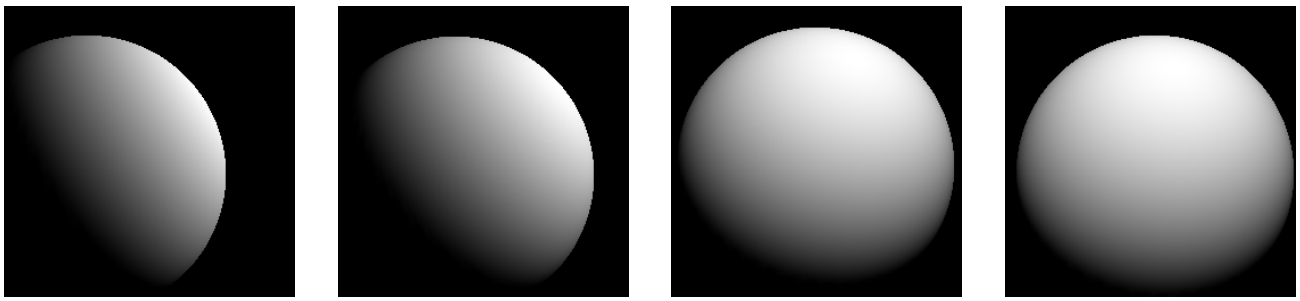

(a) $\nabla\bar{f}(x,y) = \langle 0.7474, -0.5397\rangle, \theta^* = 21^\circ,\ \theta = 22.7904^\circ$

(b) $\nabla\bar{f}(x,y) = \langle 0.6482, -0.5760\rangle, \theta^* = 37^\circ,\ \theta = 29.8680^\circ$

(c) $\nabla\bar{f}(x,y) = \langle 0.2327, -0.6067\rangle, \theta^* = 49^\circ,\ \theta = 49.4763^\circ$

(d) $\nabla\bar{f}(x,y) = \langle 0.1161, -0.5979\rangle, \theta^* = 53^\circ,\ \theta = 52.4761^\circ$

Figure 6: Lambertian spheres IV.

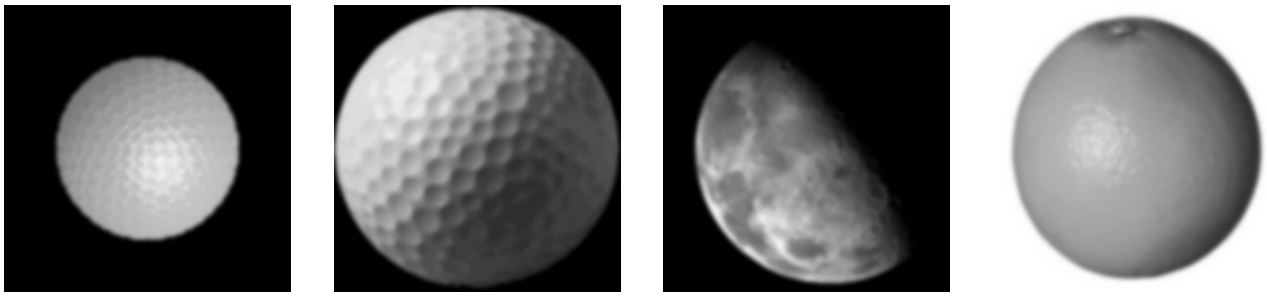

(a) $\nabla\bar{f}(x,y) = \langle 0.0735, 0.0875\rangle, \theta = 83.4406^\circ$

(b) $\nabla\bar{f}(x,y) = \langle -0.0243, -0.0320\rangle, \theta = 87.6946^\circ$

(c) $\nabla\bar{f}(x,y) = \langle -0.2437, 0.2147\rangle, \theta = 71.0466^\circ$

(d) $\nabla\bar{f}(x,y) = \langle -0.2562, -0.0006\rangle, \theta = 75.1553^\circ$

Figure 7: Textured spheres I.

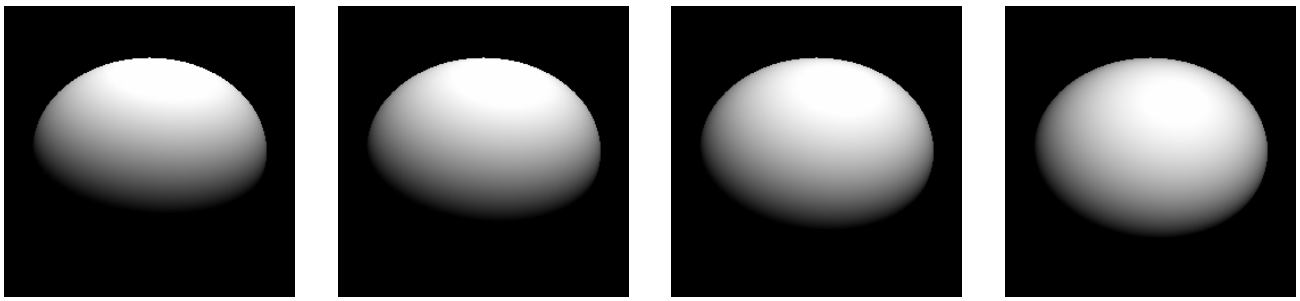

(a) $\nabla\bar{f}(x,y) = \langle 0.1494, -0.7506\rangle, \theta^* = 44^\circ,\ \theta = 40.0647^\circ$

(b) $\nabla\bar{f}(x,y) = \langle 0.1687, -0.6758\rangle, \theta^* = 50^\circ,\ \theta = 45.8524^\circ$

(c) $\nabla\bar{f}(x,y) = \langle 0.1849, -0.5652\rangle, \theta^* = 58^\circ,\ \theta = 53.5071^\circ$

(d) $\nabla\bar{f}(x,y) = \langle 0.1957, -0.4109\rangle, \theta^* = 66^\circ,\ \theta = 62.9290^\circ$

Figure 8: Lambertian ellipsoids I.

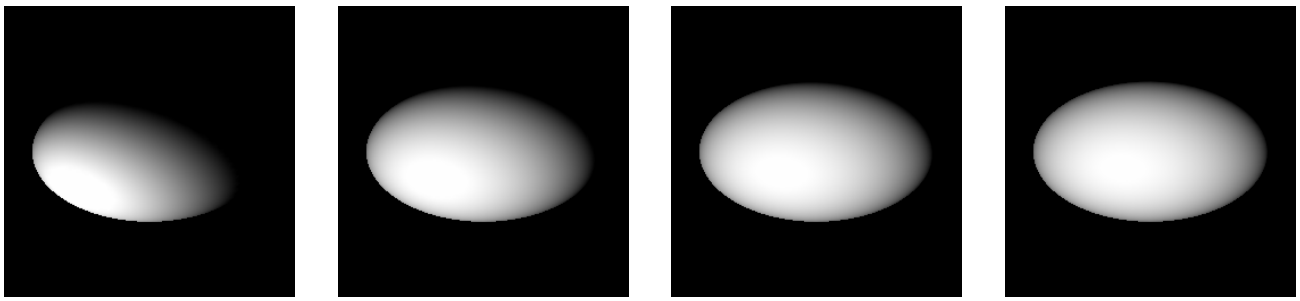

(a) $\nabla\bar{f}(x,y) = \langle -0.5917, 0.4078\rangle, \theta^* = 48^\circ,\ \theta = 44.0637^\circ$

(b) $\nabla\bar{f}(x,y) = \langle -0.3728, 0.2906\rangle, \theta^* = 66^\circ,\ \theta = 61.7877^\circ$

(c) $\nabla\bar{f}(x,y) = \langle -0.2626, 0.2247\rangle, \theta^* = 73^\circ,\ \theta = 69.7784^\circ$

(d) $\nabla\bar{f}(x,y) = \langle -0.2007, 0.1820\rangle, \theta^* = 77^\circ,\ \theta = 74.2785^\circ$

Figure 9: Lambertian ellipsoids II.

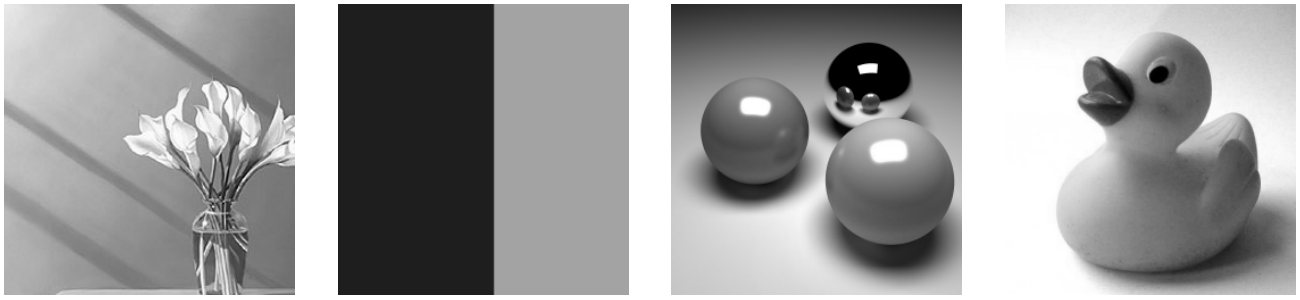

(a) $\nabla\bar{f}(x,y) = \langle -0.1886, -0.1115\rangle, \theta = 77.3433^\circ$

(b) $\nabla\bar{f}(x,y) = \langle 1.0000, 0.0000\rangle, \theta = 0.0000^\circ$

(c) $\nabla\bar{f}(x,y) = \langle -0.0040, 0.3087\rangle, \theta = 72.0168^\circ$

(d) $\nabla\bar{f}(x,y) = \langle -0.1568, 0.0683\rangle, \theta = 80.1520^\circ$

Figure 10: Natural Scenes I.