# Shape-from-shading on a cloudy day

M. S. Langer and S. W. Zucker

*McGill Research Center for Intelligent Machines, 3480 University Street, Montreal H3A2A7, Canada*

Received June 22, 1992; revision received March 30, 1993; accepted July 9, 1993

In the classical shape-from-shading model, surface luminance depends primarily on the unit surface normal. However, under diffuse lighting conditions, such as the sky on a cloudy day, luminance depends primarily on the amount of sky that is visible from each surface element, with surface normal of secondary importance. This claim is formalized in terms of a dominating sky principle and a surface aperture function. An approximately functional constraint between surface luminance and aperture emerges. It is shown how to use this constraint to recover a depth map from an image efficiently. A curious difference from the classical shape-from-shading problem is uncovered. When one assumes a point light source, the local geometric constraints of the shape-from-shading problem lie along the surface. However, in the diffuse-lighting problem, the local geometric constraints are found in a visibility field, which is defined in the free space above the surface.

## 1. INTRODUCTION

Shape-from-shading has been defined as the problem of computing a depth map from a piecewise smooth intensity image.[1] Classical solutions to this problem are based on an assumption that surface luminance is determined by the surface normal. For outdoor scenes on a sunny day, when there is a well-defined light-source direction, this classical assumption is usually valid. However, in many other situations the assumption fails. In particular, when a scene is illuminated by a diffuse light source such as the sky on a cloudy day, surface luminance is not determined by the surface normal. This is the situation with which we are concerned.

To illustrate, consider a scene composed of a single convex object resting on a ground plane. It has been argued[2] that under diffuse lighting conditions classical shape-from-shading techniques could be used to recover the shape of the object. The idea was that the region of the sky that directly illuminates a point on the convex object would be entirely determined by the surface normal at that point.

Observe, however, that these techniques would fail to recover the shape of the ground plane. Because the surface normal of the ground plane is constant, the classical shading model requires that the luminance of the ground plane also be constant. However, this requirement is clearly not met, because different points on the ground plane are exposed to different amounts of the diffuse light source. (In particular, points that are closer to the object receive more light from the diffuse source than points that are farther from the object.)

A second example should clarify this limitation of the classical model. Figure 1 shows an artificial image formed by overlapping a set of rectangles. Each rectangle has a ramp intensity in the vertical direction. Most people perceive this image as depicting a scene composed of parallel flat surfaces (e.g., a skyline or a graveyard). The percept is valid because, under diffuse lighting conditions the luminance of such protruding surfaces would vary with height. Points that are higher along a given patch would receive more illumination from the diffuse source than points below.

Clearly the percept of parallel flat surface patches is inconsistent with classical shape-from-shading, which requires that smooth luminance variations are entirely due to surface normal variations. In the skyline interpretation, the surface normals are constant along each surface patch, yet the luminance is clearly not constant.

The classical shape-from-shading model cannot, in general, account for surface shading under diffuse lighting conditions. In this paper we present a new shape-from-shading model that is designed specifically for these conditions. The main idea of the model is that under diffuse lighting conditions surface luminance is determined primarily by the amount of the diffuse source to which a surface element is exposed. We show how this model can be used to recover a depth map from a given image.

## 2. OVERVIEW

For a general scene the relationship among surface material, surface geometry, and surface luminance is governed by the radiosity equation,[3,4] which is an integral equation that explicitly relates the luminance of each surface point in a scene to the luminance of all the other surface points. In Section 3 we discuss the radiosity equation for a scene composed of Lambertian surfaces with constant albedo and illuminated by a uniform diffuse light source.

In Section 4 we present a dominating sky principle. This is an inequality that relates the luminance of surfaces in a scene to the brightness of the uniform diffuse light source. In Section 5 we introduce a useful geometric quantity, the surface aperture, that measures the solid angle of the light source that is visible from the surface. A second inequality emerges from this definition. Together these two inequalities allow us to approximate the radiosity equation by expressing the surface luminance in terms of the surface aperture and the surface albedo.

In Section 6 we extend our discussion from the geometry of incident light to the geometry of ambient light by

0740-3232/94/020467-12$06.00 

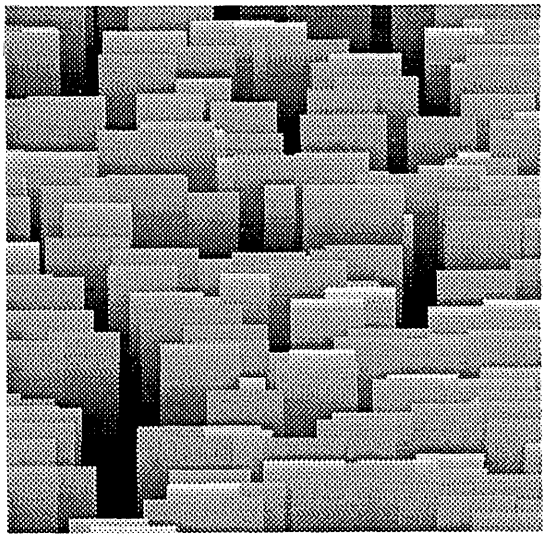

Fig. 1. Image formed by overlapping a set of rectangles, each having a vertical intensity gradient. Most people perceive a set of flat surface patches protruding from a ground plane. A skyline and a graveyard are two typical interpretations. These percepts cannot be explained by classical shape-from-shading models, which require that smooth variations in image intensity depict smooth variations in surface normal. Classical models would interpret the tombstones as curved surfaces, such as cylinders!

introducing geometric constraints that relate how much of the light source is visible from various points in the free space above the surface. In Section 7 we show how these visibility constraints can be used to compute shape-from-shading. Finally, in Section 8 we present a number of experimental results.

## 3. RADIOSITY EQUATION

As we discussed in Section 1, classical shape-from-shading models are often inappropriate for scenes illuminated by a diffuse light source. In this section we develop a framework from which a more appropriate model can be derived.

Let the free space in a scene be denoted as a point set, $\mathcal{F}$. Light is both absorbed and reflected at the surfaces in the scene, that is, at the boundary of free space, $\partial\mathcal{F}$. By definition, $\partial\mathcal{F}$ does not include points at infinity. Assume that the surfaces are smooth, so that the unit surface normal is well defined:

$$\mathbf{N}:\partial\mathcal{F} \rightarrow \text{unit sphere}.$$

Thus, for each $\mathbf{x} \in \partial\mathcal{F}$, there is a tangent plane at $\mathbf{x}$ as well as a hemisphere of directions from which $\mathbf{x}$ can receive illumination. Let $\mathcal{H}(\mathbf{x})$ denote this unit hemisphere of illuminant directions:

$$\mathcal{H}(\mathbf{x}) \equiv \{\mathbf{L}:\mathbf{L} \in \text{unit sphere}, \quad \text{and} \quad \mathbf{L} \cdot \mathbf{N}(\mathbf{x}) > 0\},$$

so that $\mathcal{H}(\mathbf{x})$ is an open subset of the unit sphere (see Fig. 2).

*Definition 1.* For any surface point $\mathbf{x} \in \partial\mathcal{F}$, the surface visibility at $\mathbf{x}$ is the set of incident directions $\nu(\mathbf{x})$ in which the diffuse light source is visible from $\mathbf{x}$. Thus the directions in which other surfaces in the scene are visible from $\mathbf{x}$ is $\nu_c(x) = \mathcal{H}(\mathbf{x})\backslash\nu(\mathbf{x})$.

Notice that for any $\mathbf{x} \in \partial\mathcal{F}$, the set $\nu(\mathbf{x})$ is an open subset of $\mathcal{H}(\mathbf{x})$. Intuitively, if the light source were visible from $x$ in direction $\mathbf{L}$, then there would exist a neighborhood of directions (containing $\mathbf{L}$) in which the source was also visible from $\mathbf{x}$.

The topology of $\nu(\mathbf{x})$ could be complicated. For example, if there were objects such as balloons or birds floating in space, then isolated regions of the sky would be occluded, and $\nu(\mathbf{x})$ would have holes. Another example is that certain objects, such as the branches of a tree, could disconnect $\nu(\mathbf{x})$ into a finite number of connected components.

It is essential to distinguish the light that is incident upon and reflected from a surface element. Let $I_{\text{in}}(\mathbf{x}, \mathbf{L})$ denote the brightness of the light ray that is incident at $\mathbf{x}$ from direction $\mathbf{L}$. The units of brightness are lumens/(steradian/square meter).[5] Let $I_{\text{out}}(\mathbf{x})$ denote the luminance of $\mathbf{x}$ also with the units lumens/(steradian/square meter).[5] Because we are assuming that the surfaces are Lambertian, $I_{\text{out}}(\mathbf{x})$ does not depend on the direction of reflection.

Let $d\Omega$ denote a small solid angle, that is, a small element of area on the hemisphere $\mathcal{H}(\mathbf{x})$. This solid angle is centered at direction $\mathbf{L} \in \mathcal{H}(\mathbf{x})$. Assuming that the albedo, $\rho$, is constant throughout the scene, we integrate over all incident directions to obtain

$$I_{\text{out}}(\mathbf{x}) = \frac{\rho}{\pi}\int_{\mathcal{H}(\mathbf{x})} I_{\text{in}}(\mathbf{x}, \mathbf{L})\mathbf{L} \cdot \mathbf{N}(\mathbf{x})d\Omega. \tag{1}$$

Next, assume that the diffuse light source has uniform brightness so that, for any $\mathbf{L} \in \nu(\mathbf{x})$,

$$I_{\text{in}}(\mathbf{x}, \mathbf{L}) \equiv I_{\mathcal{D}}.$$

Here we are ignoring nonuniformities that often arise in nature. For example, on overcast days there is a brightness peak in the zenith direction,[7] and on days when the cloud layer is thin there is a peak in the direction of the Sun.

The uniform-sky assumption permits a decomposition of Eq. (1) as follows:

$$I_{\text{out}}(\mathbf{x}) = \frac{\rho}{\pi} I_{\mathcal{D}} \int_{\nu(\mathbf{x})} \mathbf{L} \cdot \mathbf{N}(\mathbf{x})d\Omega + \frac{\rho}{\pi}\int_{\nu_c(\mathbf{x})} I_{\text{in}}(\mathbf{x}, \mathbf{L})\mathbf{L} \cdot \mathbf{N}(\mathbf{x})d\Omega. \tag{2}$$

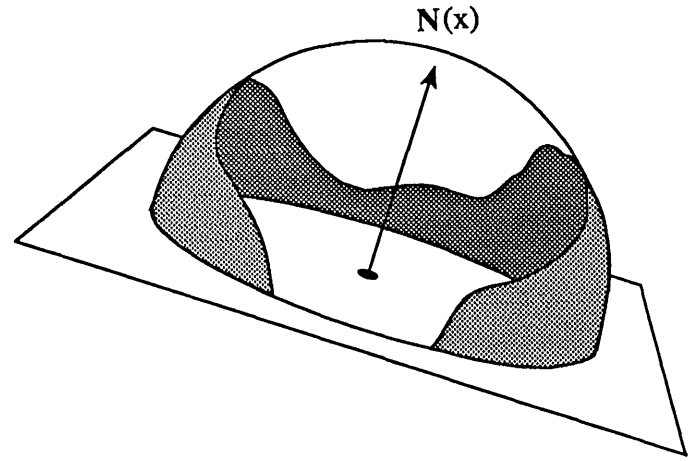


Fig. 2. Hemisphere of incident directions above a point on a surface, which can be partitioned into two sets. $\nu(\mathbf{x})$, shown by the unshaded region, denotes the set of directions in which the light source is visible from $\mathbf{x}$. Its complement, $\nu_c(x)$, denotes the set of directions pointing to other surfaces in the scene.

The first component is due to the direct illumination from the light source. The second component is due to the light reflected off other surfaces in the scene.

For each $\mathbf{L} \in \nu_c(\mathbf{x})$, let $\Pi(\mathbf{x}, \mathbf{L})$ denote the surface point seen from surface element $\mathbf{x}$ in direction $\mathbf{L}$. Thus $I_{\text{out}}[\Pi(\mathbf{x}, \mathbf{L})]$ is the luminance of the point $\Pi(\mathbf{x}, \mathbf{L})$. Neglect any light scattering within the free space so that, for any $\mathbf{x} \in \partial\mathcal{F}$ and for any $\mathbf{L} \in \nu_c(\mathbf{x})$,

$$I_{\text{in}}(\mathbf{x}, \mathbf{L}) = I_{\text{out}}[\Pi(\mathbf{x}, \mathbf{L})]. \tag{3}$$

The above assumptions can be summarized by the following model, the radiosity equation:

$$I_{\text{out}}(\mathbf{x}) = \frac{\rho}{\pi} I_{\mathcal{D}} \int_{\nu(\mathbf{x})} \mathbf{L} \cdot \mathbf{N}(\mathbf{x}) d\Omega + \frac{\rho}{\pi} \int_{\nu_c(\mathbf{x})} I_{\text{out}}[\Pi(\mathbf{x}, \mathbf{L})] \mathbf{L} \cdot \mathbf{N}(\mathbf{x}) d\Omega. \tag{4}$$

Three shading factors are explicitly represented in this equation. Mutual shadowing is represented by the partition of the integral into two domains. Mutual illumination is represented by the nonzero luminances within the second integral. The classical dependence on surface normal also plays a role.

Notice that if $\nu(\mathbf{x})$ were independent of $\mathbf{x}$, then the above model would reduce to the classical model, in which an equivalent point source is well defined. As we have argued, however, under diffuse lighting conditions, $\nu(\mathbf{x})$ often depends quite strongly on $\mathbf{x}$.

Computational techniques for solving the radiosity equation (under far more general conditions) have been developed within the computer graphics community (e.g., Ref. 8). Their goal was to render accurately a scene with known geometry and known surface material. Surface geometry is approximated as a mesh of $N$ planar surface elements, and the radiosity equation is approximated as a set of $N$ linear equations.

A related technique has been applied to the classical shape-from-shading problem.[9] The strategy was to estimate surface shape with classical techniques (that is, ignoring mutual illumination and mutual shadowing) and then to adjust the computed surface iteratively until its luminance was consistent with the original image and satisfied the radiosity equation. It was argued that this method would converge, provided that the initial estimate of the surface was sufficiently accurate, that is, provided that the classical shading effects dominated.

In the present paper we develop a technique for estimating shape-from-shading under diffuse lighting conditions. Our image formation model is quite different from the classical model in that we assume luminance is determined primarily by shadowing rather than by the surface normal. In Sections 4 and 5 we show how to obtain this model by an approximation of the radiosity equation.

## 4. DOMINATING SKY PRINCIPLE

Our discussion of shading under diffuse lighting conditions begins with a few familiar observations. When an object such as a tree or a bird is viewed against the sky, the silhouette of the object is clearly demarcated, although the luminance variations within the silhouette may be difficult to discern. Similarly, at dawn or dusk many objects appear only as silhouettes. The reason for this phenomenon is that under diffuse lighting conditions the sky is typically much brighter than the ground. In particular the following can be shown (see Appendix A for a proof):

*Lemma 1.* Dominating sky principle: Let a scene be illuminated by a diffuse light source with uniform brightness $I_{\mathcal{D}}$, and let the surfaces be Lambertian with constant albedo $\rho$. Then the luminance of the surface is bounded as follows:

$$0 \le I_{\text{out}}(\mathbf{x}) \le \rho I_{\mathcal{D}}.$$

In one large compilation of natural scenes,[10] the mean albedo was found to be approximately 0.15.[11] This datum is consistent with a separate study in which measurements of ambient light were made over a large number of natural habitats.[12] The data from the latter study indicate that on overcast days the brightness of the sky is often 10 times as great as the brightness of the ground.

Many insect species provide wonderful support for the dominating sky principle in the form of a mechanism that has evolved for detecting the horizon.[13] It has been shown that certain insects cannot stabilize their roll (the rotation about the longitudinal axis of the body) when they are placed in a uniformly lit environment, that is, one that does not contain a horizon. Such insects are believed to stabilize their flight by visually controlling the attitudes of their bodies with respect to the visible sky, thus maximizing the difference in illumination above and below them.

## 5. SURFACE APERTURE AND MUTUAL SHADOWING

The dominating sky principle implies that there is a relationship between the luminance of a surface element and the amount of the diffuse light source that is visible from that element.[14] If the entire source were visible, then the luminance would be maximal, whereas if none of the source were visible, then the luminance would be relatively low. A simple example should clarify this idea.

Figure 3(a) shows a scene consisting of a cylinder lying next to a deep gully. The brightest region of the scene is the top of the cylinder, because from there the entire diffuse source is visible. The luminance decreases continuously as one moves to the bottom of the cylinder. From there none of the diffuse source is visible. Similarly, as one descends into the gully, the diffuse source gradually becomes occluded, and the surface becomes darker. This shadowing effect, that different regions of the diffuse source are visible from different points on the surfaces, is represented geometrically in Fig. 3(b).

In this section we address the question of how much of the luminance variation along a surface is due to this shadowing effect. The following definition allows us to specify the degree of shadowing:

*Definition 2.* The surface aperture at $\mathbf{x} \in \partial\mathcal{F}$ is the percentage of incident directions in which the hemispheric source is visible. Formally,

$$A(\mathbf{x}) \equiv \frac{1}{2\pi} \int_{\nu(\mathbf{x})} d\Omega.$$

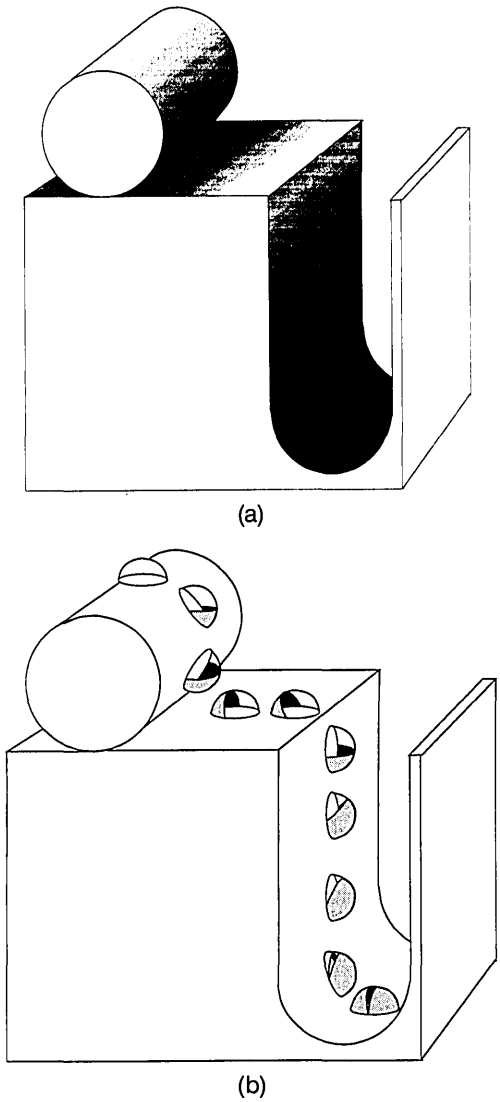


Fig. 3. (a) Surface luminance function. (b) Shadowing effects represented geometrically by the partition of the hemisphere of incident directions for each point. Notice that points in which less of the sky is visible tend to have lower luminance.

For the scene in Fig. 3, the aperture would be 1 along the top of the cylinder and 0 along the bottom of the cylinder. The aperture would increase as one moved along the ground plane and away from the cylinder and would decrease as one descended along the wall of the gully.

A useful bound on surface aperture is the following:

*Lemma 2.*

$$A(\mathbf{x}) \le \frac{1}{\pi}\int_{\nu(\mathbf{x})} \mathbf{L}\cdot\mathbf{N}(\mathbf{x})\mathrm{d}\Omega \le A(\mathbf{x})[2 - A(\mathbf{x})]. \quad (5)$$

These bounds are calculated in Appendix B. The idea is to fix $A(\mathbf{x})$ and to consider what $\nu(\mathbf{x})$ would have to be to maximize or minimize the given integral. For a maximum, $\nu(\mathbf{x})$ would be a set of directions bounded by a cone centered on $\mathbf{N}(\mathbf{x})$. For the minimum, $\nu(\mathbf{x})$ would be an annulus of directions crowded around the horizon of $\mathcal{H}(\mathbf{x})$.

The radiosity equation and the dominating sky principle may be combined to yield

$$\frac{\rho}{\pi}I_{\mathcal{D}}\int_{\nu(\mathbf{x})}\mathbf{L}\cdot\mathbf{N}(\mathbf{x})\mathrm{d}\Omega \le I_{\mathrm{out}}(\mathbf{x}) \le \rho^2 I_{\mathcal{D}} + (1-\rho)\frac{\rho}{\pi}I_{\mathcal{D}}\int_{\nu(\mathbf{x})}\mathbf{L}\cdot\mathbf{N}(\mathbf{x})\mathrm{d}\Omega. \quad (6)$$

Substituting inequality (5) into inequality (6) yields

$$\rho I_{\mathcal{D}} A(\mathbf{x})^2 \le I_{\mathrm{out}}(\mathbf{x}) \le \rho^2 I_{\mathcal{D}} + (1-\rho)\rho I_{\mathcal{D}} A(\mathbf{x})[2 - A(\mathbf{x})], \quad (7)$$

which can be rewritten as follows:

*Proposition 1.* Aperture–luminance inequality:

$$\max\left(0, 1 - \left\{\frac{1}{1-\rho}\left[1 - \frac{1_{\mathrm{out}}(\mathbf{x})}{\rho I_{\mathcal{D}}}\right]\right\}^{1/2}\right) \le A(\mathbf{x}) \le \left[\frac{I_{\mathrm{out}}(\mathbf{x})}{\rho I_{\mathcal{D}}}\right]^{1/2}. \quad (8)$$

How tight are these bounds of relation (8) in practice? To answer this question, we examine the values of $I_{\mathrm{out}}(\mathbf{x})$ and $A(\mathbf{x})$ for a number of surfaces. Figure 4(a) is a mesh plot of a drapery surface whose aperture values are well distributed over the range of 0 to 1. This surface was rendered with 40 point sources (an approximation of a uniform diffuse light source) by the use of the radiosity equation. (Our rendering algorithm differs from the one described in Ref. 7. We plan to present details in a future paper.)

Notice that the bounds of inequality (8) are quite loose relative to the plotted values of $I_{\mathrm{out}}(\mathbf{x})$ and $A(\mathbf{x})$. In particular the data points are clustered near the middle of the permitted range. This clustering was expected. For a surface point to attain the upper (or lower) bound of inequality (8), it is necessary that the upper (or lower) bounds of inequalities (5) and (6) be attained simultaneously. This latter requirement is rarely met for two reasons. First, inequalities (5) and (6) represent distinct constraints, so an attainment of the bounds of one relation implies nothing about the attainment of the bounds of the other. Second, as is discussed in Appendixes A and B, the conditions under which each of the bounds of inequalities (5) and (6) are attained are rarely met.

To carry this investigation one step further, we examine a well-known situation that is problematic for classical shape from shading. Concave surfaces that have large albedos typically do not satisfy the classical shading model, because the mutual illumination effect for such surfaces is large.[15,16] Does a similar difficulty arise under uniform hemispheric lighting conditions?

We address this question by rendering a smooth concave surface, again using the radiosity equation. The depth map is a $50 \times 50$ concave Gaussian function having a standard deviation of 15 pixels and a depth ranging from 0 to 50 pixels. In Figs. 5(a)–5(c), rendered images are shown for three different values of $\rho$. In Figs. 5(d)–5(f), scatter plots of the $[I_{\mathrm{out}}(\mathbf{x}), A(\mathbf{x})]$ pairs are shown.

Clearly the albedo affects the mean value of $A(\mathbf{x})$ for a given value of $I_{\mathrm{out}}(\mathbf{x})$. That is, when $\rho$ is large, the values of $I_{\mathrm{out}}(\mathbf{x})$ are skewed toward larger luminances. However, the variance of $A(\mathbf{x})$ for a given value of $I_{\mathrm{out}}(\mathbf{x})$ is roughly the same for all three albedos. The same qualitative observation can be made from the scatter plots of Fig. 4.

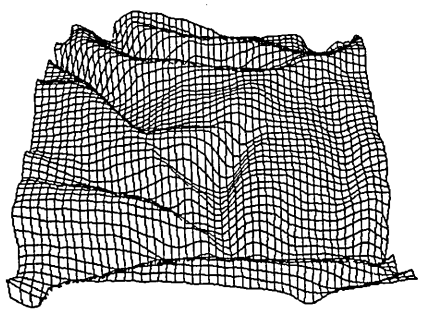

(a)

(b)

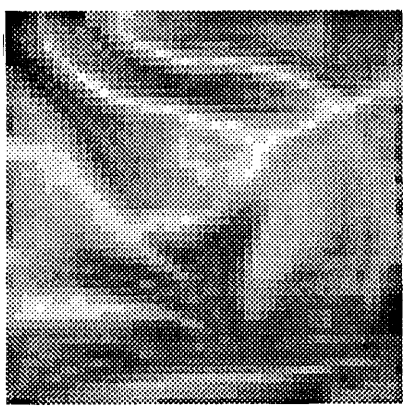

(c)

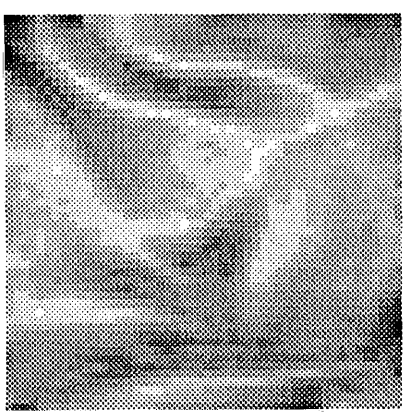

(d)

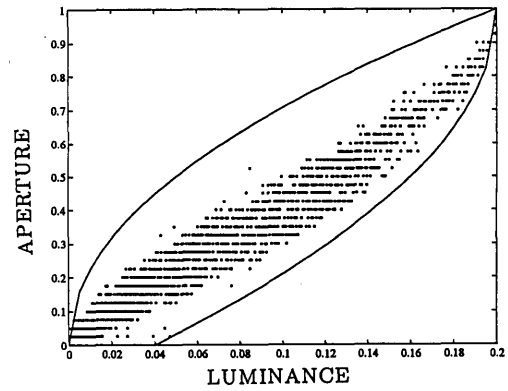


(e)

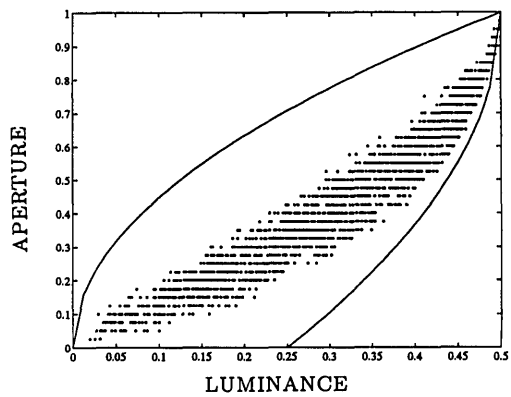


(f)

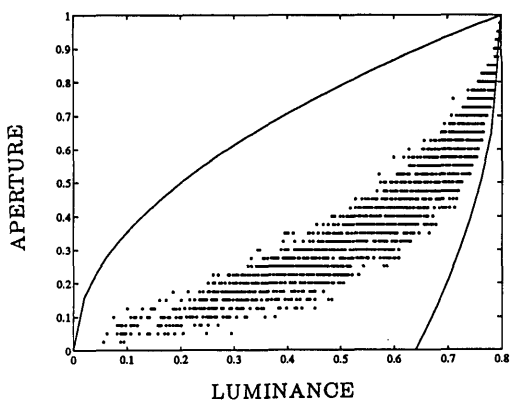


(g)

Fig. 4. (a) 50 × 50 smooth surface. The depth values of the surface span a range of 0 to 25. (b), (c), (d) Surface rendered according to the radiosity equation for three values of the albedo (0.2, 0.5, 0.8, respectively). The uniform hemispheric source was approximated with 40 point sources. (e)–(g) Scatter plot of aperture versus luminance for each of the three albedos. The two solid curves represent the upper and lower bounds of the aperture–luminance inequality. These bounds are quite weak relative to the values found for a typical surface.

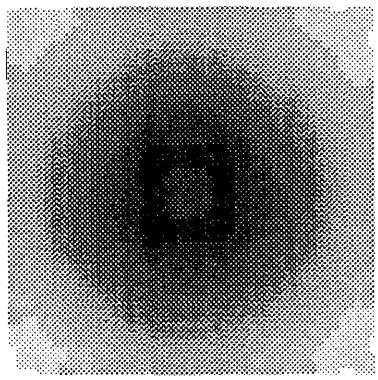

(a)

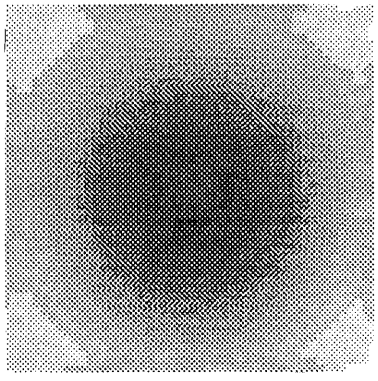

(b)

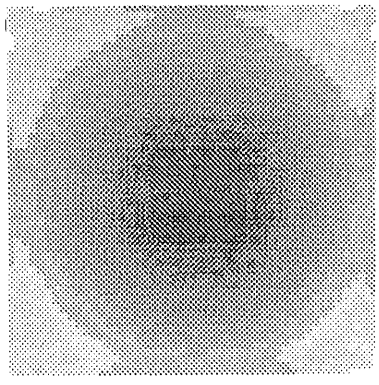

(c)

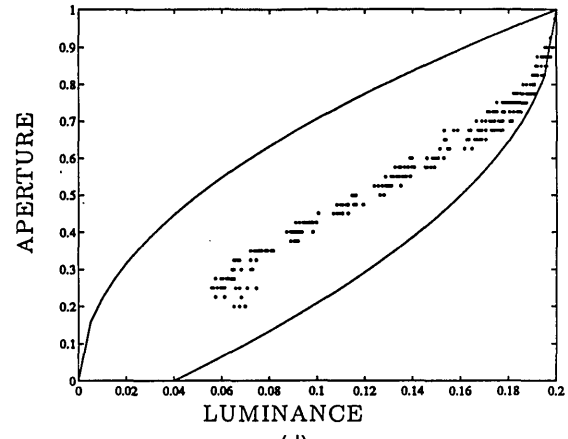


(d)

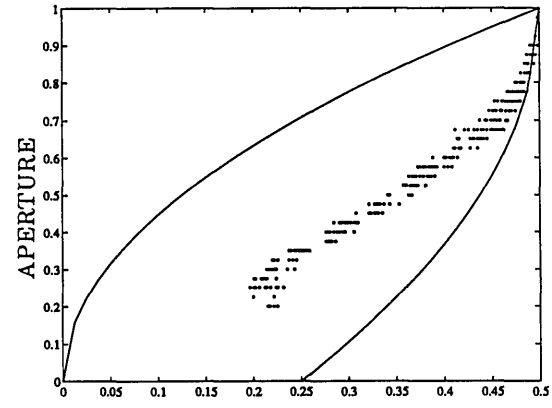


(e)

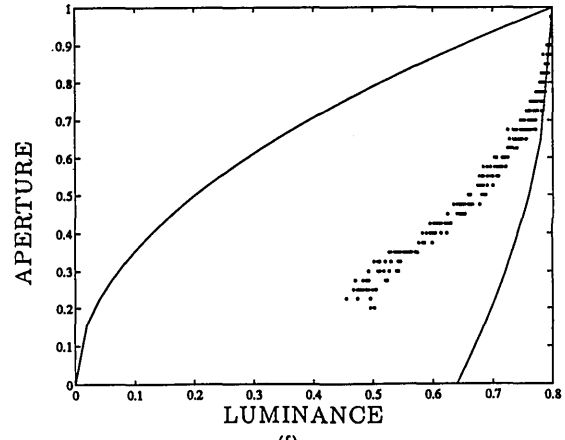


(f)

Fig. 5. (a), (b), (c) Concave 50 × 50 Gaussian function having standard deviation 15 pixels and a maximum depth of 50, rendered according to the radiosity equation with three values of $\rho$ (0.2, 0.5, 0.8, respectively). (d)–(f) Scatter plots of aperture versus luminance. (The leakage of the data beyond the bounds is a quantization effect.) Clearly, the average aperture for a given luminance varies with $\rho$. However, the variation of the aperture for a given luminance is roughly the same for all three values of $\rho$. This suggests that the variation in $A(\mathbf{x})$ for a given value of $I(\mathbf{x})$ is independent of $\rho$.

The data suggest that the effect of mutual illumination is to skew the surface luminance distribution. In particular, as the albedo of a surface is increased, the distribution is skewed toward larger values. This result should be familiar to anyone who has skied. On a cloudy day, the snow-covered hills (which have an albedo close to 1) have little contrast because the luminance distribution is greatly skewed toward a maximal value. As a result the hills appear flat. This issue is discussed again in Subsection 8.A, in which various choices of albedo are considered for a given image.

The above experiments suggest that one can estimate surface aperture by considering the average of the bounds of relation (8). Let $\tilde{A}(\mathbf{x})$ denote this estimate, so that

$$\tilde{A}(\mathbf{x}) \equiv \frac{1}{2}\left[\left[\frac{I_{\text{out}}(\mathbf{x})}{\rho I_{\mathcal{D}}}\right]^{1/2} + \max\left(0, 1 - \left\{\frac{1}{1-\rho}\left[1 - \frac{I_{\text{out}}(\mathbf{x})}{\rho I_{\mathcal{D}}}\right]\right\}^{1/2}\right)\right]. \quad (9)$$

To summarize, under uniform diffuse lighting conditions the luminance of a surface depends quite strongly on the surface aperture, that is, on the mutual shadowing of surfaces. This dependence is quite different from classical shape-from-shading models, in which it is assumed that surface luminance is determined primarily by the surface normal.

## 6. VISIBILITY FIELD

Thus far we have developed three functions: visibility, $\nu(\mathbf{x})$; luminance, $I_{\text{out}}(\mathbf{x})$; and aperture, $A(\mathbf{x})$. Each of these was defined along the surfaces in a scene. We now extend the domains of $\nu(\mathbf{x})$ and $A(\mathbf{x})$ to the free space, $\mathcal{F}$. That is, we extend our discussion from the geometry of incident light to the geometry of ambient light. This extension is useful because there are strong local constraints on the geometry of ambient light.

*Definition 3.* For any $\mathbf{x} \in \mathcal{F}$, let $\nu(\mathbf{x})$ be the set of unit directions in which the sky is visible from $\mathbf{x}$. We refer to $\nu(\mathbf{x})$ as the visibility field at $\mathbf{x}$.

There are strong local constraints on the visibility field. Intuitively, for any $\mathbf{x} \in \mathcal{F}$, if the diffuse light source were visible in a particular direction $\mathbf{L}$, then any translation within $\mathcal{F}$ and parallel to $\mathbf{L}$ would not affect this visibility. This constraint is formally expressed as follows:

*Local visibility constraint.* Let $\mathbf{x} \in \mathcal{F}$ and $\epsilon > 0$ be such that the open ball centered at $\mathbf{x}$ and of radius $\epsilon$ is entirely contained in $\mathcal{F}$. For any direction $\mathbf{L}$ and for any $\delta \in [-\epsilon, \epsilon]$, we have $\mathbf{L} \in \nu(\mathbf{x})$ if and only if $\mathbf{L} \in \nu(\mathbf{x} + \delta\mathbf{L})$.

The local visibility constraint allows us to compute shape from shading. The main concept (which is developed in Section 7) is that, if the visibility field were known in a neighborhood within free space, then the local visibility constraint could be used to expand locally this neighborhood.

The visibility field is related to a classical construction in illumination engineering.[5] The light field is the vector field on $\mathcal{F}$ defined by

$$\mathbf{D}(\mathbf{x}) \equiv -\int_{\nu(\mathbf{x})} I_{\mathcal{D}} \mathbf{L} \, d\Omega.$$

This vector field, $\mathbf{D}(\mathbf{x})$, is physically meaningful. It is the flux density of radiant energy coming directly from the diffuse light source. Although we do not make use of the light field when computing shape-from-shading, we make use of the following scalar field:

*Definition 4.* For any $\mathbf{x} \in \mathcal{F}$, the aperture field at $\mathbf{x}$ is the percentage of the hemispheric source that is visible from $\mathbf{x}$:

$$A(\mathbf{x}) \equiv \frac{1}{2\pi}\int_{\nu(\mathbf{x})} d\Omega.$$

Both the visibility and aperture fields play an important role in our shape-from-shading computation. It is this computation to which we now turn.

## 7. COMPUTATIONAL PROBLEM AND SOLUTION

A viewer is faced with the following computational problem. Given an image, compute a depth map that is consistent with the image, where consistency is defined with respect to Eq. (9). There are two stages to the solution. The first is to transform the image into an estimate of the surface aperture function. The second stage is to use this estimated aperture function to compute the depth map of a surface.

### A. Definition of Quantized Variables

Because digital images are used as input, it is necessary to quantize the variables of the problem. We use the superscript * to remind the reader that a variable has been quantized.

Let $\mathcal{Z}$ be the set of integers. We begin by quantizing space with a cubic lattice of nodes:

$$\mathcal{Z}^3 \equiv \{\mathbf{x} = (x, y, n): x, y, n \in \mathcal{Z}\}.$$

The set of free nodes in a scene is denoted $\mathcal{F}^*$. These are the nodes lying strictly above the surfaces. The surface nodes are then denoted $\partial\mathcal{F}^*$.

Assume that the image occupies a small solid angle of view, so that perspective distortions can be ignored. That is, the image projection is orthographic. Let $(0, 0, 1)$ be the viewing direction. $\mathcal{P}$ is an $N \times N$ lattice of pixels, so that a depth map is a function:

$$z: \mathcal{P} \rightarrow \{0, 1, \ldots\}.$$

A viewed surface is thus a set of nodes:

$$\partial\mathcal{F}^* \equiv \{[x, y, z(x, y)]: (x, y) \in \mathcal{P}\} \subset \mathcal{Z}^3.$$

Because of the limited field of view, it is necessary to make an assumption about the image boundary conditions. Assume that beyond the image window the surface has constant depth. Intuitively, the viewer is looking directly down at a pit that has been excavated from the ground. Further assume that the boundary depth is identical to the depth of the node nearest to the viewer.

The quantized hemisphere of possible light source directions is denoted as

$$\mathcal{H}^* \equiv \{\mathbf{L}^k: k = 1, \ldots M\}.$$

These are the $M$ directions above the ground plane from which a surface node can potentially be illuminated. $\mathcal{H}^*$ is defined such that each direction is a unit vector pointing from the origin to a node in $\mathcal{Z}^3$. Figure 6 depicts two quantizations of $\mathcal{H}^*$ with $M = 32$ and 64.

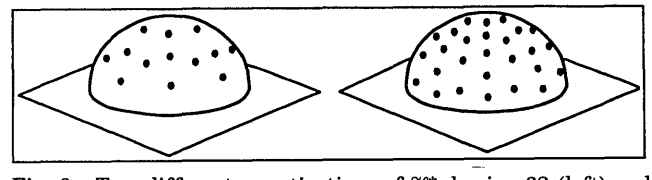

Fig. 6. Two different quantizations of $\mathcal{H}^*$, having 32 (left) and 64 (right) light-source directions, respectively. Each of these unit vectors points from the origin to a node in $\mathcal{Z}^3$.

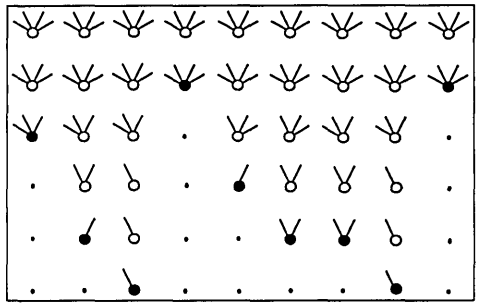

Fig. 7. Example of a quantized visibility field for five light-source directions. Free nodes, $\mathcal{F}^*$, are represented by open disks. Surface nodes, $\partial\mathcal{F}$, are represented by filled disks. Nodes below the surface are represented as points.

*Definition 5.* For any node $\mathbf{x} \in \mathcal{F}^* \cup \partial\mathcal{F}^*$, the quantized visibility field at $\mathbf{x}$, $\nu^*(\mathbf{x}) \subseteq \mathcal{H}^*$, is the finite set of directions in which the sky is visible from $\mathbf{x}$. (See, for example, Fig. 7.) For any finite set $U$, the number of elements of $U$ is denoted $|U|$.

*Definition 6.* For any node $\mathbf{x} \in \mathcal{F}^* \cup \partial\mathcal{F}^*$, the quantized aperture field is

$$A^*(\mathbf{x}) \equiv \frac{|\nu^*(\mathbf{x})|}{|\mathcal{H}^*|}.$$

Finally, for any node $\mathbf{x} = (x, y, n)$, denote

$$\nu_n^*(x, y) \equiv \nu^*(\mathbf{x}), \qquad A_n^*(x, y) \equiv A^*(\mathbf{x}).$$

### B. From Image Brightness to Surface Aperture

Let $I^*(x, y)$ be the quantized image brightness, with units of lumens/(steradian/square meter). In particular the maximum image brightness is denoted $I_{\max}{}^*$. According to the assumed image boundary conditions, the brightest surface node has an aperture equal to 1. Hence the following estimate is obtained from Eq. (9):

$$\tilde{A}^*(\mathbf{x}) \equiv \frac{1}{2}\left[\left[\frac{I_{\text{out}}(\mathbf{x})}{I_{\max}{}^*}\right]^{1/2} + \max\left(0, 1 - \left\{\frac{1}{1-\rho}\left[1 - \frac{I_{\text{out}}(\mathbf{x})}{I_{\max}{}^*}\right]\right\}^{1/2}\right)\right]. \quad (10)$$

Notice that it is necessary to choose a particular value of $\rho$. This choice should be familiar to the reader, as it is also required in the classical shape-from-shading problem. As we illustrate in Subsection 8.A, although the amplitude of the recovered surface varies with the error between the actual $\rho$ and the chosen $\rho$, the location of the hills and valleys of the surface is approximately preserved.

### C. From Surface Aperture to Depth

The second stage of the computation is to recover a depth map from the estimated surface aperture function. The algorithm is iterative, beginning at the depth of the node nearest to the viewer. The main idea is that, if the visibility field were known up to depth $n$, then it could be extended to depth $n + 1$ with the following analog of the local visibility constraints discussed in Section 6.

*Quantized local visibility constraints.* Let $\mathbf{x} \in \mathcal{F}^* \cup \partial\mathcal{F}^*$ and $\mathbf{L} \in \mathcal{H}^*$. If $\mathbf{x}' \in \mathcal{Z}^3$ is the nearest node to $\mathbf{x}$ in direction $\mathbf{L}$, then $\mathbf{L} \in \nu^*(\mathbf{x})$ if and only if all the following hold:

1. $\mathbf{x}' \in \mathcal{F}^*$, that is, $\mathbf{x}'$ lies strictly above the surface;
2. $\mathbf{L} \in \nu^*(\mathbf{x}')$;
3. any node on the open line segment between $\mathbf{x}$ and $\mathbf{x}'$ belongs to $\mathcal{F}^*$.

It is crucial to appreciate that these are local constraints. For each node at fixed depth, the directions in which the diffuse source is visible from that node can be computed by examination of a small neighborhood above that node. The computation begins at the depth of the node nearest the viewer. For all nodes at this depth, the entire source is assumed to be visible. Then, for any given depth $n$, and for any node $(x, y, n) \in \mathcal{F}^* \cup \partial\mathcal{F}^*$, the visibility field at $(x, y, n)$ is computed with the quantized local visibility constraints. If the aperture of $(x, y, n)$ were larger than the estimated aperture value of the corresponding pixel, then it would follow that $(x, y, n) \in \mathcal{F}^*$. For pixel $(x, y)$, the surface is reached when the computed aperture has decreased to the estimated aperture.

*Algorithm.* Given an estimated aperture function, $\tilde{A}^*$, compute a depth map, $z$:

$\forall(x, y) \in \mathcal{P}, \quad z(x, y) := 0;$
$\forall(x, y) \in \mathcal{P}, \quad \forall n \le 0, \nu_n^*(x, y) := \mathcal{H}^*;$
$n := 0;$
REPEAT
$\forall(x, y) \in \mathcal{P},$
IF $z(x, y) = n$
  compute $\nu_n^*(x, y)$ with the local visibility
  constraints; compute $A_n^*(x, y)$ with relation (6);
IF $A_n^*(x, y) > \tilde{A}^*(x, y)$
  $z(x, y) := z(x, y) + 1;$
$n := n + 1;$
UNTIL $\forall(x, y) \in \mathcal{P}, \quad A_n^*(x, y) \le \tilde{A}^*(x, y)$ ■

We have proved a number of results concerning the behavior of the algorithm.[17]

*Proposition 2.* The algorithm always converges.

*Proof.* Because of the image boundary conditions, all nodes beyond a certain depth, $z_{\max}$, would have minimum aperture. For points beyond this depth, the quantized diffuse source would not be visible. It follows that the algorithm could not proceed beyond $z_{\max}$ iterations. ■

*Proposition 3.* Suppose that for a given an aperture function $A^*$ there exists a depth map that is consistent with $A^*$. Then, the algorithm computes the shallowest depth map that is consistent with $A^*$.

To summarize this section, we have partitioned the problem of computing shape-from-shading into two problems. First, a surface aperture function is estimated from a given image by use of an image formation model based entirely on mutual shadowing. Second, a depth map of a surface having that estimated aperture function is computed.

## 8. EXPERIMENTAL RESULTS

We address three questions concerning the performance of our algorithm.

### A. Choice of Albedo

In the first step of the algorithm, the surface aperture function is estimated. This estimate depends on a chosen

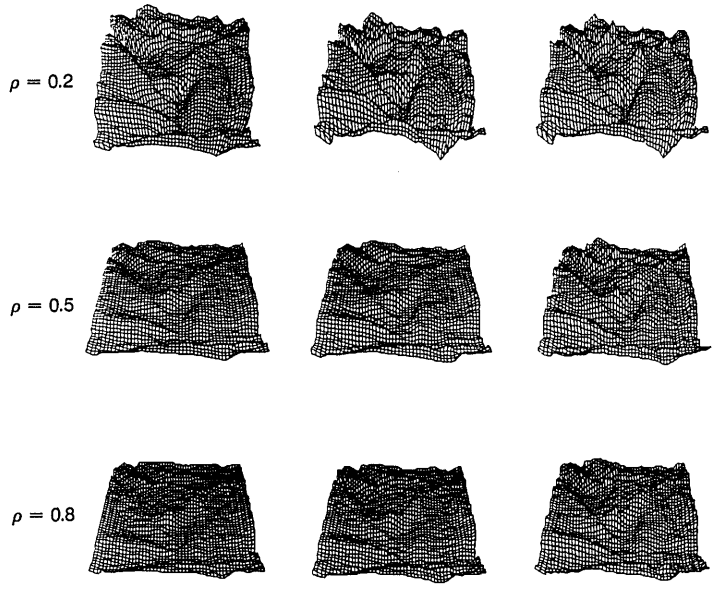


Fig. 8. Shape-from-shading on a cloudy day algorithm applied to each of the three rendered images of Fig. 4. Each column corresponds to a single rendered image of Fig. 4. Each row above corresponds to a different chosen value of $\rho$ (0.2, 0.5, 0.8). The three surfaces along the diagonal thus correspond to the correctly chosen values of $\rho$. See the text for a discussion of the errors.

**Table 1. Mean Squared Depth Error of the Computed Surfaces of Fig. 8[a]**

| $\rho$ | 0.2 | 0.5 | 0.8 |
|---|---|---|---|
| 0.2 | 64.6 | 16.5 | 30.0 |
| 0.5 | 99.0 | 8.8 | 22.5 |
| 0.8 | 144.8 | 7.0 | 10.2 |

[a]In general one would expect the errors to be smallest along the diagonal (when the albedo is chosen correctly). Of course, because inequality (6) is an estimate, an incorrectly chosen albedo may produce a smaller error. Indeed, for the example shown, the diagonal elements do not always have the smallest error.

albedo, $\rho$. The following experiment illustrates how the computed surface varies with the chosen $\rho$.

Consider again the depth map of Fig. 4 and the rendered images of Figs. 4(b)–4(d). We apply our shape-from-shading algorithm to each of these images, using three different albedos for each image. The resulting nine depth maps are shown in Fig. 8, and the mean squared errors between the original depth map and the computed depth maps are shown in Table 1. In general one would expect that the errors are smallest along the diagonal, that is, when the albedo is chosen correctly. Of course, because Eq. (9) is defined as an estimate of $\tilde{A}(\mathbf{x})$, cases may arise in which an incorrectly chosen albedo produces a smaller error. Indeed, for the example shown, the diagonal elements do not always have the smallest errors.

When the chosen albedo is larger than the actual albedo, the reconstructed depth map has greater amplitude of variation than the actual depth map. The opposite effect occurs when the chosen albedo is too small. To appreciate why this occurs, recall how the rendered images of Figs. 4 and 5 varied with the actual albedo. As the albedo is increased, the surface luminance becomes skewed toward a maximal value. Thus, if the albedo were overestimated, then the image brightnesses would be less skewed than one would expect. This wider distribution of brightnesses would then be accounted for by a greater amplitude of variation in the computed depth map. Similarly, if the albedo were underestimated, then the skewing in the actual image brightnesses would be larger than expected. (That is, the image contrast would be smaller than expected.) Again, this unexpected brightness distribution would be accounted for by the computed depth map. It may be possible to use the distribution of image brightnesses in choosing the albedo of a given image. We are currently investigating this idea.

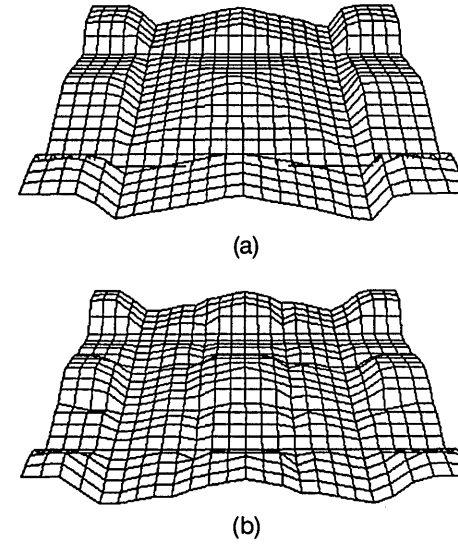


Fig. 9. Because $z$ and $\mathcal{H}^*$ are quantized, the two algorithms—aperture from depth (see Appendix C) and depth from aperture—are not inverses of each other. When these two algorithms are run in succession on an original constant depth map, the resulting computed depth is not constant. (a) Mesh plot of a computed depth map for an original 50 × 50 depth map having constant depth ($z \equiv 40$) and with relatively coarse quantization parameters ($M = 32$, and $z$ quantized to unit steps). (b) Mesh plot for a finer quantization ($M = 64$, and $z$ quantized to half-unit steps).

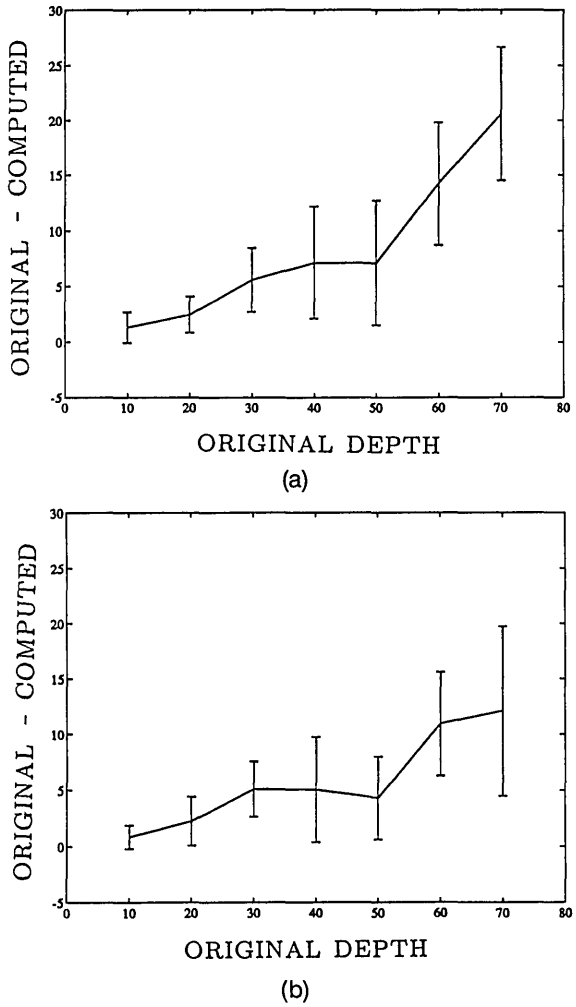


Fig. 10. Aperture from depth and depth from aperture algorithms were run in succession on a set of depth maps, each having constant depth. The mean and the standard deviation of the difference between the original and the computed depth map are plotted. (a) Coarse quantization ($M = 32$, and $z$ quantized to unit steps), (b) fine quantization ($M = 64$, with $\hat{z}$ quantized to half-unit steps). See the text for a discussion of the errors.

### B. Quantization Errors

Although our mathematical results hold in principle, quantization effects are an issue and require careful investigation. In the next experiment, seven constant depth maps were created. For each depth map, the aperture function was computed (see Appendix C). Then, a depth map was computed from each of these aperture functions. If the two computations (aperture from depth and depth from aperture) were inverses of each other, then their successive application would reproduce the original constant depth maps. However, this is not the case. These errors in the computed depth map are entirely due to quantization.

Figures 9(a) and 9(b) show the computed depth maps corresponding to $z \equiv 40$ for coarse and fine quantization, respectively. The symmetry of the error surface reflects the quantization of $\mathcal{H}^*$.

In Figs. 10(a) and 10(b) the mean and the standard deviation of the difference between the original constant depth map and the computed depth map are plotted for both coarse and fine quantization. Notice first that the errors are smaller for the finer quantization, although they are not nearly fine enough to be ignored. Second, the errors are quite biased. The computed depth map is much shallower than the original. (The bias illustrates the effect mentioned in proposition 3.) Third, the errors are greater for deeper depth maps. The reason is that as the surfaces become deeper, the aperture becomes small. In this case the diffuse source effectively becomes a local source, such as a window or a lamp. Clearly, the quantization of $\mathcal{H}(\mathbf{x})$ is no longer appropriate in this limiting case. This issue is discussed in greater depth in Ref. 17.

### C. Relation between Depth and Surface Luminance

Finally, we use our theory to address a classical issue in shading from the artistic literature. As Leonardo da Vinci observed, "among bodies equal in size and distance, that which shines the more brightly seems to the eye nearer."[18] Notice, for example, how in the rendered images of Fig. 4 our own visual systems are not able to infer precisely the depth map of the drapery surface, although there is a strong sense that the bright regions are hills and the dark regions are valleys. Artists often take advantage of this darker-is-deeper heuristic when rendering drapery surfaces.[19]

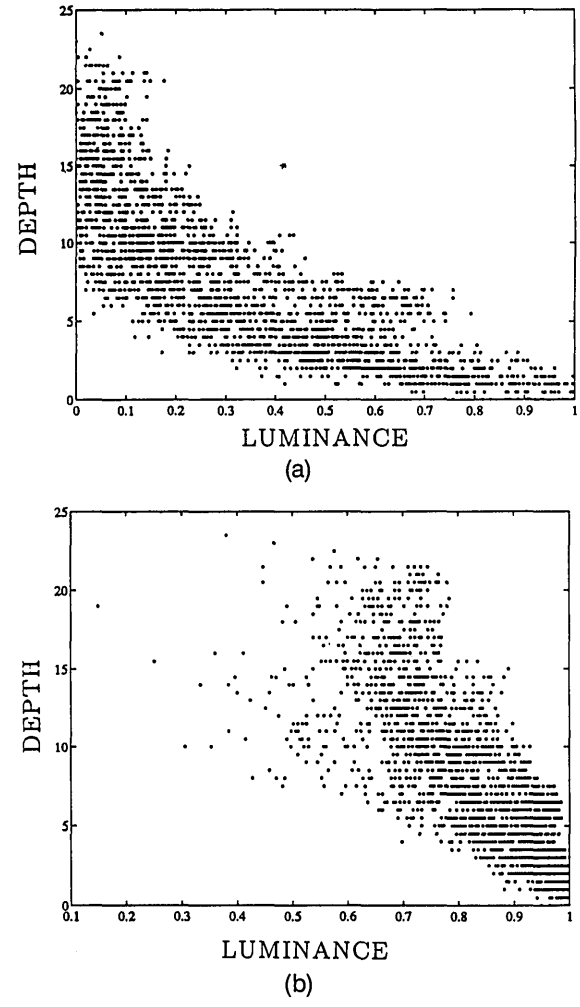


Fig. 11. Scatter plot of luminance–depth pairs for the rendered images of Fig. 4: (a) $\rho = 0.2$, (b) $\rho = 0.8$. A darker-is-deeper heuristic is valid only in a loose statistical sense. The large variance of the relation implies that the heuristic is unreliable for recovering a depth map from an image.

The shape-from-shading algorithm developed in this paper also (at first glance) suggests a heuristic that darker surface nodes tend to be deeper. To examine the validity of this heuristic, we again consider the smooth depth map and rendered images of Fig. 4. Figure 11 shows a scatter plot of luminance–depth pairs for these images. The plots suggest that, on average, darker surface elements do have greater depth. However, the large variance of this relation implies that the heuristic is unreliable.

A few examples should clarify why the heuristic can fail. Wrinkles on a person's face or cracks in the floor are dark but not deep. Here, darkness results from small surface aperture, rather than from large depth. Conversely, a surface such as a large valley could have great depth, but it need not be dark. In particular, if the width of a valley were much greater than the depth, then the luminance of the valley could be nearly as great as that of the surrounding hills. The conclusion is that darkness is neither necessary nor sufficient for deepness but rather is coupled to the surface aperture function.

## 9. DISCUSSION

There are many types of shape-from-shading problems. The most basic is to distinguish the sky from the ground. We have discussed how a dominating sky principle, which holds that the sky is much brighter than the ground in natural environments, could allow a flying organism to orient itself. A simple visual system would be sufficient for such a task.

A slightly more complex visual system could distinguish different brightnesses coming from the surfaces in a scene. An organism might behave differentially toward these different brightnesses. In particular, if an organism tended to move toward relatively low (or high) brightnesses, then it would on average move into (or out of) holes or shadows. Whether this behavior was advantageous to the organism would depend, of course, on the particular organism and its environment.

A more complex visual system would be able to use the variations in brightness to compute the geometric structure of surfaces in a scene. There are many different lighting conditions under which an organism might solve this problem. The classical example of a point light source has been studied in great detail. Another example is a scene illuminated by a nearby light source, such as a lamp or fire. For such a scene, a model relating surface luminance and scene geometry would have to take account of how the lighting varies with the position of a surface element in the scene.

In this paper we have considered scenes that are illuminated by a uniform diffuse light source, such as the sky on a cloudy day. We have developed, from first principles, a method for computing a depth map of a Lambertian surface from an image. The first half of the paper is an argument that surface luminance depends primarily on the surface aperture function and that the surface normal is a secondary factor.

In the second half of the paper, we discuss local visibility constraints on the geometry of ambient light. An algorithm for computing a depth map is derived from these constraints. The algorithm uses simple operations at each computational node and local connectivity between nodes. Conceptually, the geometric constraints that allow one to compute shape from shading are along the surface for the classical point source situation but above the surface for a diffuse light source. It is nevertheless pleasing that the latter, seemingly more complex, situation still admits of efficient solutions.

## APPENDIX A

We begin by addressing a folk belief held by many that, by a fortuitous arrangement of surfaces, it could happen that a surface element is brighter than the uniform, hemispheric source. The following lemma shows this belief to be unfounded. From the lemma, we derive the dominating sky principle.

*Lemma 3.* $\forall \mathbf{x} \in \partial\mathcal{F}, \quad I_{\text{out}}(\mathbf{x}) \le I_{\mathcal{D}}$.

*Proof.* Let $\mathbf{x}^*$ be the brightest surface point in the scene, so that $I_{\text{out}}(\mathbf{x}) \le I_{\text{out}}(\mathbf{x}^*)$ for any $\mathbf{x} \in \partial\mathcal{F}$. Suppose that $I_{\mathcal{D}} < I_{\text{out}}(\mathbf{x}^*)$ (we derive a contradiction). From the radiosity equation,

$$\begin{aligned} I_{\text{out}}(\mathbf{x}^*) &= \frac{\rho}{\pi}\int_{\nu(\mathbf{x}^*)} I_{\mathcal{D}}\mathbf{L}\cdot\mathbf{N}(\mathbf{x}^*)\,d\Omega \\ &\quad + \frac{\rho}{\pi}\int_{\nu_c(\mathbf{x}^*)} I_{\text{out}}(\mathbf{x}^*,\mathbf{L})\mathbf{L}\cdot\mathbf{N}(\mathbf{x})\,d\Omega \\ &< \frac{\rho}{\pi}\int_{\mathcal{H}(\mathbf{x}^*)} I_{\text{out}}(\mathbf{x}^*)\mathbf{L}\cdot\mathbf{N}(\mathbf{x}^*)\,d\Omega\,. \end{aligned}$$

Thus $I_{\text{out}}(\mathbf{x}^*) < \rho I_{\text{out}}(\mathbf{x}^*)$, which is a contradiction because $0 \le \rho \le 1$. ■

From the radiosity equation and lemma 3 we have

$$\begin{aligned} I_{\text{out}}(\mathbf{x}) &= \frac{\rho}{\pi}\int_{\nu(\mathbf{x})} I_{\mathcal{D}}\mathbf{L}\cdot\mathbf{N}(\mathbf{x})\,d\Omega \\ &\quad + \frac{\rho}{\pi}\int_{\nu_c(\mathbf{x})} I_{\text{out}}[\Pi(\mathbf{x},\mathbf{L})]\mathbf{L}\cdot\mathbf{N}(\mathbf{x})\,d\Omega \\ &\le \frac{\rho}{\pi}\int_{\mathcal{H}(\mathbf{x})} I_{\mathcal{D}}\mathbf{L}\cdot\mathbf{N}(\mathbf{x})\,d\Omega \\ &= \rho I_{\mathcal{D}}\,, \end{aligned}$$

which proves the dominating sky principle.

It is important to observe that the bounds of the dominating sky principle are attained only under exceptional circumstances. For any $\mathbf{x} \in \mathcal{F}$, the upper bound would be attained only if all surfaces visible from $\mathbf{x}$ had luminance equal to $I_{\mathcal{D}}$. However, if $\rho < 1$, then (by the dominating sky principle) this last condition would be met only in the trivial case that no surfaces were visible from $\mathbf{x}$.

The lower bound of the dominating sky principle would be attained in the limiting case that none of the diffuse source was visible from $\mathbf{x}$ and that those surfaces that were visible from $\mathbf{x}$ all had zero luminance.

## APPENDIX B

*Proof of lemma 2.* Given $\mathbf{N}(\mathbf{x})$ and $A(\mathbf{x})$, define a spherical coordinate parameterization of $\mathcal{H}(\mathbf{x})$ with a pole at $\mathbf{N}(\mathbf{x})$ as follows:

$$\mathbf{L}(\theta,\phi):[0,2\pi][0,\pi/2] \rightarrow \text{unit sphere},$$

where

$$\mathbf{L}(\theta,\phi)\cdot\mathbf{N}(\mathbf{x}) = \cos\phi,$$
$$d\Omega = \sin\phi d\phi d\theta.$$

Let $\phi_{\min}$ and $\phi_{\max}$ be the unique real numbers such that

$$A(\mathbf{x}) = \int_0^{2\pi}\int_0^{\phi_{\max}} \sin\phi d\phi d\theta$$
$$= \int_0^{2\pi}\int_{\phi_{\min}}^{\pi/2} \sin\phi d\phi d\theta.$$

A straightforward calculation shows that

$$A(\mathbf{x}) = 1 - \cos\phi_{\max} = \cos\phi_{\min}. \quad \text{(B1)}$$

Define

$$\nu_{\max} \equiv \{\mathbf{L}(\theta,\phi): 0 \le \theta \le 2\pi, 0 < \phi < \phi_{\max}\},$$
$$\nu_{\min} \equiv \{\mathbf{L}(\theta,\phi): 0 \le \theta \le 2\pi, \phi_{\min} < \phi < \pi/2\}.$$

These sets of directions are the result of taking any given solid angle of directions and, first, crowding them around $\mathbf{N}(\mathbf{x})$ to form $\nu_{\max}$ and, second, crowding them into an annulus around the horizon to form $\nu_{\min}$. These definitions imply the following:

$$\frac{1}{\pi}\int_{\nu_{\min}} \mathbf{L}\cdot\mathbf{N}(\mathbf{x})d\Omega \le \frac{1}{\pi}\int_{\nu(\mathbf{x})} \mathbf{L}\cdot\mathbf{N}(\mathbf{x})d\Omega$$
$$\le \frac{1}{\pi}\int_{\nu_{\max}} \mathbf{L}\cdot\mathbf{N}(\mathbf{x})d\Omega. \quad \text{(B2)}$$

This inequality does not make any topological assumptions about the connectedness or number of holes of $\nu(\mathbf{x})$, beyond requiring that the integrals be defined.

A straightforward calculation shows that

$$\frac{1}{\pi}\int_{\nu_{\min}} \mathbf{L}\cdot\mathbf{N}(\mathbf{x})d\Omega = \cos^2\phi_{\min},$$
$$\frac{1}{\pi}\int_{\nu_{\max}} \mathbf{L}\cdot\mathbf{N}(\mathbf{x})d\Omega = 1 - \cos^2\phi_{\max}.$$

Using Eq. (B1), we substitute the above equations into Eq. (B2), yielding

$$A(\mathbf{x})^2 \le \frac{1}{\pi}\int_{\nu(\mathbf{x})} \mathbf{L}\cdot\mathbf{N}(\mathbf{x})d\Omega \le A(\mathbf{x})[2 - A(\mathbf{x})]. \quad \blacksquare$$

Observe that $\nu_{\max}$ and $\nu_{\min}$ are quite symmetric. Because such symmetry rarely arises in natural scenes, it follows that the bounds of lemma 2 are rarely attained.

## APPENDIX C

*Algorithm*: Given a depth map $z$, compute the surface aperture, $A^*$.
$\forall(x,y) \in \mathcal{P}, \forall n \le 0, \nu_n^*(x,y) := \mathcal{H}^*$;
$n := 0$;
REPEAT
  $\forall(x,y) \in \mathcal{P}$,
  IF $z(x,y) = n$
    compute $\nu_n^*(x,y)$, using the local visibility constraints;
  $n := n + 1$;
UNTIL $\forall(x,y) \in \mathcal{P}, n > z(x,y)$;
$\forall(x,y) \in \mathcal{P}$, compute $A^*(x,y)$ with relation (6).

## REFERENCES


1. B. K. P. Horn and M. J. Brooks, eds., *Shape From Shading* (MIT, Cambridge, Mass., 1989).
2. B. K. P. Horn and R. W. Sjoberg, "Calculating the reflectance map," Appl. Opt. **18,** 1770–1779 (1979).
3. E. M. Sparrow and R. D. Cess, *Radiation Heat Transfer* (McGraw-Hill, New York, 1978).
4. J. J. Koenderink and A. J. van Doorn, "Geometrical modes as a general method to treat diffuse interreflections in radiometry," J. Opt. Soc. Am. **73,** 843–850 (1983).
5. A. Gershun, "The light field," J. Math. Phys. **18,** 51–151 (1939).
6. M. V. Klein and T. E. Furtak, *Optics* (Wiley, New York, 1986).
7. J. J. Koenderink and W. A. Richards, "Why is snow so bright?" J. Opt. Soc. Am. A **9,** 643–648 (1992).
8. C. M. Goral, K. E. Torrence, D. P. Greenberg, and B. Battaile, "Modelling the interaction of light between diffuse surfaces," Comput. Graphics **18,** 213–222 (1984).
9. S. K. Nayar, K. Ickeuchi, and T. Kanade, "Shape from interreflections," Int. J. Comput. Vision **6,** 173–195 (1991).
10. E. L. Krinov, *Spectral Reflectance Properties of Natural Materials,* NRC Technical Translation 439 (National Research Council of Canada, Ottawa, 1971).
11. W. A. Richards, "Lightness scale from intensity distributions," Appl. Opt. **21,** 2569–2582 (1982).
12. J. P. Hailman, "Environmental light and conspicuous colors," in *The Bahavioral Significance of Color,* E. H. Burtt, ed. (Garland, New York, 1979).
13. R. Wehner, "Spatial vision in arthropods," in *Handbook of Sensory Physiology,* H. Autrum, ed., Vol. C of Comparative Physiology and Evolution of Vision in Invertebrates (Springer-Verlag, Berlin, 1981).
14. I. M. S. Langer and S. W. Zucker, "Qualitative shape from active shading," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* (Institute of Electrical and Electronics Engineers, New York, 1992), pp. 713–715.
15. B. K. P. Horn, "Understanding image intensities," Artif. Intell. **8,** 201–231 (1977).
16. D. Forsyth and A. Zisserman, "Reflections on shading," IEEE Trans. Pattern Anal. Mach. Intell. **13,** 671–679 (1991).
17. I. M. S. Langer, "The computational geometry of light," Ph.D. dissertation (McGill University, Montreal, Quebec, Canada, 1994).
18. E. MacCurdy, ed., *The Notebooks of Leonardo da Vinci* (Jonathan Cape, London, 1938), p. 332.
19. K. Nicolaides, *The Natural Way to Draw* (Houghton Mifflin, Boston, Mass., 1941).