

Casting Light on Illumination: A Computational Model and Dimensional Analysis of Sources

MICHAEL S. LANGER*

NEC Research Institute, 4 Independence Way, Princeton, New Jersey 08540

AND

STEVEN W. ZUCKER

Computer Science Department, Yale University, New Haven, Connecticut 06520-8285

Received November 15, 1995; accepted November 20, 1996

We attempt to cast light on illumination by introducing a computational model of light sources and illumination. Spatially varying illumination is computed using coordinate transformations on the manifold of rays in space, and the ray manifold is shown to be 4 dimensional. The illumination model can be used to recover shape from shading in situations when direct illumination varies across space and provides the basis for relating different types of light sources (e.g., an overcast sky, the sun, a panel light, and a spotlight). A formal definition for light sources emerges that generalizes the standard definition of a source as an emitter and leads to a 4-D light source hypercube in which the different types of sources may be embedded. © 1997 Academic Press

1. INTRODUCTION

The visual appearance of a scene depends on a variety of physical factors: the shape and material of object surfaces, the 3-D spatial layout of surfaces, and the illumination. Each factor plays a different role in visual function: Surface shape and material support object recognition, 3-D spatial layout supports navigation, and illumination supports inferences about which parts of a scene are visible (e.g., dark shadows typically hide the objects within them).

Interestingly, humans seem to be exquisitely sensitive to shape or material, especially in complex scenes, but are less sensitive to illumination. For example, we notice sudden changes in reflectance properties of skin, due to blushing, sweating, or paleness, as well as changes in shape, such as when we see a person smiling or frowning. In contrast, we often fail to notice small changes in illumina-

tion, such as when a thin cloud passes in front of the sun and the illumination becomes more diffuse. Similarly, we often have difficulty guessing the lighting configuration in a scene, such as in a magazine photograph, even though we may feel confident of the surface shapes, materials, and 3-D layout. Based on these observations, it is tempting to believe that we judge illumination poorly relative to how well we judge shape, material, and 3-D layout and that in many cases, we cannot judge illumination at all.

Yet, from a computational perspective, it would be puzzling if the visual system were less sensitive to illumination variations than to variations of shape, material, and 3-D layout. All of these variables are intrinsically coupled in the physics of image formation, and it is impossible to say in general that one variable determines scene appearance more or less than another. To analyze image data, one would expect that all these variables need to be considered. To put it another way, if vision were a chain of inferences from image properties to scene properties, then shouldn't the chain be only as strong as its weakest link?

A common approach to this weakest link puzzle is to claim there are "photometric invariants" in the visual world [9, 15, 14, 8, 27] that allow visual systems to recover geometry and material without having to explicitly recover the illumination. While we agree that certain invariants do exist and are useful, we argue that they are insufficient. For example, when a dark region in an image is inferred to be a shadow rather than dark colored material, this inference involves illumination. When a floor is illuminated by a proximal source such as a lamp, and surface points beneath the lamp are observed to be brighter than points far from the lamp, the inference about what is occurring involves illumination. The same holds when a surface brightens as it curves toward a light source and the shape is inferred. We believe that illumination is not an invariant

* To whom correspondence should be addressed. E-mail: langer@research.nj.nec.com.

of our visual inferences, but rather plays a central role in them, just as shape, material, and 3-D layout do. Our goal in this paper is to provide a theory that suggests this belief is plausible.

To build our intuition toward this end, observe that illumination is related to the way that we function within space. It is straightforward that shadows disturb reading and that lamps must be hidden from view to avoid glare. Candles and spotlights are used for different reasons, and lasers are hardly ever used except as pointers. More subtle, however, is how lighting is designed to introduce overall properties of space. In a restaurant, for example, lighting might be arranged to divide large spaces into small cones for personal intimacy. In a shop, lighting is designed to draw attention to display areas and to direct paths of customer circulation [37]. Stadium lighting is designed to make a space appear uniform with regard to paths of moving people and objects. Thus there is a relationship between illumination and space, and this relationship will play a central role in our approach.

In this paper, we present a model of how illumination varies across space and specify data structures and algorithms for computing spatially varying illumination. The model is at an abstract level: it specifies what is computed and how. We *are* motivated by biology in that we want to understand the visual problems that biological systems must be solving, but we emphasize that we do not intend the model to be taken literally as it relates to mechanisms of the brain. This distinction between an abstract level of description and the biological mechanisms has a long history in computer vision; see Marr [23]. We need not repeat those arguments here, but merely remind the reader to keep this distinction in mind.

2. PREVIOUS WORK

Before we begin, let us review the models that are currently used in computer vision and discuss their limitations. Examples of light source models include a point source at infinity [10], a proximal source [4], a linear source [13], a uniform hemispheric source [19], a sum of point sources at infinity [3], and various structured spotlights [32, 28]. A key limitation of this collection of models is that there exists no framework for relating them. Such a framework is important. To see why, consider the problem of recovering shape from shading. Current algorithms are designed for specific ideal light sources, such as a point source at infinity [12, 29] or uniform hemispheric source [19]. Yet, in a non-contrived situation, the *type* of source and the surface shape must be inferred simultaneously, presumably by tolerating parametric errors in each. A general algorithm would require a general framework for handling different types of light sources.

A second limitation is that, unlike the sources mentioned




FIG. 1. Two distinct free spaces are shown in light grey and are enclosed by dotted lines.

above, real light sources are nonideal. The sun is not a point, but rather subtends a finite angle and is accompanied by a blue sky whose contribution to illumination is nonnegligible. A real proximal source such as a lamp is not a point either, and it is typically nonisotropic [26]. Vision systems should expect light sources to be nonideal and to treat nonideal variations appropriately, rather than, for example, as unstructured Gaussian noise.

The only general framework for illumination modeling in vision research that we are aware of is the “plenoptic function” of Adelson and Bergen [1]. We argue that their model is not computational, however, since it fails to provide constraints that allow illumination to be computed. In particular, the plenoptic function fails to take advantage of the most fundamental constraint on illumination, which is that light travels along rays. This is the constraint on which our computational model will be based.

3. RAY MANIFOLD

Our model is based on a well-known empirical law of radiometry which states that, in the absence of atmospheric emission or scattering, and above the scale of diffraction, the intensity of light is constant along a geometric ray [33]. This law holds whether a ray is coming directly from a source, whether it is reflected from a surface, or in the latter case, whether the surface has mirror, glossy, or matte reflectance. The law is central to the physics of light. It is the reason that rays are used in geometric optics. Specifically, the law makes the plenoptic function redundant by one dimension. This redundancy provides the main constraint for our computational model.

We introduce a model of illumination that is based on the set of light rays in a scene. In this section, we define this set of rays and show how it may be parameterized. We begin by defining a *free space*, \mathcal{F} , to be a bounded, open, connected subset of \mathfrak{R}^3 , which is free of objects and has a piecewise smooth boundary, $\partial\mathcal{F}$. By an “object,” we mean anything that would scatter, reflect, emit, or absorb light. Two examples of free spaces are shown in Fig. 1. Note that the boundary of a free space, $\partial\mathcal{F}$, need not be




FIG. 2. Two rays in a non-convex free space. Are the rays neighbors or not?

restricted to the surfaces of actual physical objects. For example, a window or doorway could be part of $\partial\mathcal{F}$, as could an imaginary bubble enclosing an outdoor scene. Note also that a free space may be nonconvex. That is, it may contain holes, which are objects that reflect, emit, absorb, or scatter light.

We next define the set of rays in a given free space. For any two points $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^3$, let $(\mathbf{x}_1, \mathbf{x}_2)$ denote the open directed line segment from \mathbf{x}_1 to \mathbf{x}_2 , and let $[\mathbf{x}_1, \mathbf{x}_2]$ denote the closed directed line segment from \mathbf{x}_1 to \mathbf{x}_2 . Let $V(\mathbf{x}_1, \mathbf{x}_2)$ be a binary *visibility* function which is 1 if \mathbf{x}_1 and \mathbf{x}_2 are visible from one another, and 0 otherwise. That is, $V(\mathbf{x}_1, \mathbf{x}_2) = 1$ if and only if $(\mathbf{x}_1, \mathbf{x}_2) \subseteq \mathcal{F}$. The set of rays, $\mathcal{M}(\mathcal{F})$, in a free space is as follows.

DEFINITION 1. Given a free space, \mathcal{F} , the set of rays, $\mathcal{M}(\mathcal{F})$, is the set of all $[\mathbf{x}_1, \mathbf{x}_2]$ such that

- $\mathbf{x}_1 \neq \mathbf{x}_2$,
- $\mathbf{x}_1 \in \partial\mathcal{F}, \mathbf{x}_2 \in \partial\mathcal{F}$,
- $V(\mathbf{x}_1, \mathbf{x}_2) = 1$.

Each ray has a point of origin and a point of termination in $\partial\mathcal{F}$ and, other than these endpoints, each ray is strictly contained in \mathcal{F} . Note that the ray $[\mathbf{x}_1, \mathbf{x}_2]$ is distinct from the ray $[\mathbf{x}_2, \mathbf{x}_1]$ since these two rays have opposite points of origin and termination. Also, two rays may be collinear but still distinct; for example, consider two rays on opposite sides of an object.

Many parameterizations of the set of rays $\mathcal{M}(\mathcal{F})$ are possible. A common one used in computer graphics [7] is to parameterize rays by their endpoints. Mathematically, this defines an embedding of $\mathcal{M}(\mathcal{F})$ into $\partial\mathcal{F} \times \partial\mathcal{F}$. The indicator function for this embedding is the visibility function, $V(\mathbf{x}_1, \mathbf{x}_2)$. Note that this embedding is 1–1, but not onto. That is, it is injective but not surjective. If free space is nonconvex, then there will be pairs of points which are not visible from one another.

The embedding of $\mathcal{M}(\mathcal{F})$ into $\partial\mathcal{F} \times \partial\mathcal{F}$ obscures $\mathcal{M}(\mathcal{F})$'s topological properties. To see this, consider the two rays shown in Fig. 2 and ask, Are the rays neighbors

or not? One could argue that the rays are neighbors in that they have similar directions and similar endpoints. Alternatively, one could argue they are not neighbors in that their points of origin are so far apart. The neighborhood relation seems to depend on the chosen parameterization.

To clarify this topological issue, we introduce a novel parameterization of rays, using concepts of coordinate charts and atlases from differential geometry [34]. The coordinate charts and transformations between them will be the basis of our analysis of light sources in Section 6 and of our computational model in Section 7. For each point in free space, $\mathbf{x}_0 = (x_0, y_0, z_0) \in \mathcal{F}$, consider the three planes

$$\mathcal{P}_{x_0} \equiv \{(x, y, z) : x = x_0\} \subset \mathbb{R}^3,$$

$$\mathcal{P}_{y_0} \equiv \{(x, y, z) : y = y_0\} \subset \mathbb{R}^3,$$

$$\mathcal{P}_{z_0} \equiv \{(x, y, z) : z = z_0\} \subset \mathbb{R}^3,$$

which we call *coordinate planes*. The set of rays in $\mathcal{M}(\mathcal{F})$ that pierce the coordinate plane \mathcal{P}_{z_0} in the positive z direction is denoted $\mathcal{U}_{z_0}^+ \subset \mathcal{M}(\mathcal{F})$. Let $\mathbf{r} \in \mathcal{U}_{z_0}^+$ be the ray shown in Fig. 3. This ray passes through \mathbf{x}_0 with direction $(p_0, q_0, 1) \in \mathbb{R}^3$. We define the mapping

$$\varphi_{z_0}^+ : \mathcal{U}_{z_0}^+ \rightarrow \mathbb{R}^4,$$

such that




FIG. 3. A ray \mathbf{r} passing through (x_0, y_0, z_0) and in direction $(p_0, q_0, 1)$. This ray may be parameterized as $\varphi_{z_0}^+(\mathbf{r}) = (x_0, y_0, p_0, q_0)$.

$$\varphi_{z_0}^+(\mathbf{r}) \equiv (x_0, y_0, p_0, q_0).$$

$\varphi_{z_0}^+$ defines a *local coordinate chart* of $\mathcal{M}(\mathcal{F})$.

Two observations should be made. First, we have avoided using angles to parameterize ray directions. Instead, we parameterize rays by projecting the hemisphere of unit directions onto a plane. This is analogous to the “gradient space” parameterization of surface normals which is classical in computer vision [22]. The reason for doing this, as we will see later, is that it simplifies the notation considerably by avoiding trigonometric expressions.

Second, recalling Fig. 2 and the discussion thereof, we now see that the question of whether two rays are neighbors depends entirely on the coordinate chart one chooses. The distance between the two rays in the (x, y, p, q) space is well defined within one chart only, and this distance varies continuously from chart to chart.

To continue our development, we have shown how the rays in a single coordinate are parameterized. But how does the parameterization of a ray change from one coordinate chart to another? Let \mathcal{P}_{z_1} be a second coordinate plane through which \mathbf{r} passes. The parameterization of a ray \mathbf{r} in the two coordinate charts, $\varphi_{z_1}^+$ and $\varphi_{z_0}^+$, are related by

$$\varphi_{z_1}^+ \circ (\varphi_{z_0}^+)^{-1}(x_0, y_0, p_0, q_0) = (x_0 + (z_1 - z_0)p_0, y_0 + (z_1 - z_0)q_0, p_0, q_0). \quad (1)$$

Note that the transformation from one coordinate chart to another is linear. As we will see in Sections 6 and 7, this linearity property makes the transformations cleaner to analyze and compute. This advantage is analogous to that of the gradient space parameterization of surface normals used in [22].

Finally, we state a basic theoretical result about the set of rays in a given scene. This result is proved and discussed in the Appendix.

PROPOSITION 1 (Ray manifold). *Given a free space \mathcal{F} , the set of rays $\mathcal{M}(\mathcal{F})$ is a 4-D manifold.*

4. RADIANCE

In the last section, we examined the geometry of rays in a given free space. In this section, we examine their density. The density of rays at point \mathbf{x} and in a direction \mathbf{u} is defined as the flux of light at \mathbf{x} per unit area perpendicular to \mathbf{u} and per unit solid angle. In radiometry, this density is called *radiance*. Radiance has units $\text{W m}^{-2} \text{sr}^{-1}$. (We ignore spectral properties of radiance.) The following is a well-known law of radiometry [33] which we mentioned earlier and now state in terms of rays and radiance.

OBSERVATION 1 (law of conservation of radiance). *In*

the absence of emission and scattering, and above the scale of diffraction, radiance is a nonnegative function on the set of rays, $\mathcal{M}(\mathcal{F})$, i.e.,

$$R: \mathcal{M}(\mathcal{F}) \rightarrow [0, \infty).$$

We rewrite Observation 1 in coordinate chart form, by letting $R_{z_0}^+(x_0, y_0, p_0, q_0)$ denote the radiance of the ray that passes through $\mathbf{x} = (x_0, y_0, z_0)$ in direction $\mathbf{u} = (p_0, q_0, 1)$. If this ray also passes through the coordinate plane \mathcal{P}_{z_1} then

$$R_{z_0}^+(x_0, y_0, p_0, q_0) = R_{z_1}^+(x_0 + (z_1 - z_0)p_0, y_0 + (z_1 - z_0)q_0, p_0, q_0). \quad (2)$$

We use this constraint throughout the remainder of the paper.

A few points are worth stressing before we proceed further. First, the above law allows no emission or scattering in free space. This precludes the interior of a candle flame or cloud, where there is emission and scattering of light, respectively. Radiance is not constant within such volumes, and the law does not apply. We are assuming that a free space \mathcal{F} satisfies the above law and that volumes such as the interior of a candle flame and a cloud lie outside of free space. We do so in the same way that we require the interior an opaque object to lie outside free space.

Second, we observe that the constraint of Eq. (2) was not mentioned in the plenoptic function analysis of [1]. Rather, the plenoptic function allows an arbitrary value of radiance at each point \mathbf{x} and direction \mathbf{u} (and at each wavelength and time). The plenoptic function is thus redundant by one dimension [1]. This redundancy will be the basis for our computational model. (This point about the redundancy of the plenoptic function was also made in [21] in comparing their image-based rendering method to that of [24].)

Third, as an aside, we note that for certain computations, especially in computer graphics, a redundant representation of rays may still be useful. For example, a 6-D parameterization of rays known as Plücker coordinates is useful for efficiently computing ray intersections [35]. This technique has been used in computer graphics to compute visibility relations in densely occluded environments [39]. Another example of a useful yet redundant representation is the 5-D parameterization of $\mathcal{M}(\mathcal{F})$ which is used in the ray-tracing method of [2].

5. WHAT IS A LIGHT SOURCE?

The ray manifold, $\mathcal{M}(\mathcal{F})$, of a given scene may be partitioned into two subsets: a set of source rays, and a set of non-source rays. This is the extension of a familiar idea,

e.g., in computer graphics, of defining certain surfaces in the scene as sources and other surfaces as reflectors. The partition of the ray manifold is useful, as we will see in Section 6, because it allows us to characterize a set of source rays in terms of its radiance distribution.

Defining a partition of the ray manifold into source and non-source rays in a way that makes sense both psychologically and physically is not straightforward, however. To appreciate the subtlety of this problem, consider the example of an outdoor scene on an overcast day. As before, we consider \mathcal{F} to be a finite volume within the scene. Most agree that the sky is a source, and so the set of source rays in $\mathcal{M}(\mathcal{F})$ is just the subset which, when extended backwards, would reach the sky. The subtlety arises when we ask whether a white piece of paper in the scene is a source. Most would answer no, even though the radiance of the paper could be just as great as that of the sky. Why is the sky considered a source but a white piece of paper not?

One distinction between the sky and the piece of paper is that the rays from the sky transmit light into free space across its boundary while the rays from the piece of paper merely reflect light back into free space. This suggests that transmission of light may be sufficient for a region on the boundary of free space to be a source. But there are problems with this answer. A translucent vase is certainly not a source, even though it transmits light into free space. Similarly, a doorway to a room need not be a source, even though light is transmitted through it into the room (e.g., a doorway leading outside would be a source, but a doorway leading to a closet would not be).

To partition a given ray manifold into source and non-source rays, it is therefore not enough to know the emitted, transmitted, and reflected components of each ray. It seems we must also know the physics of what is happening beyond the boundary of free space; *i.e.*, we have to consider the sun behind the cloud and the cloud beyond the doorway. This is unsatisfying, however, for two reasons. First, since the radiance in a given free space \mathcal{F} is entirely determined by the physics at the boundary $\partial\mathcal{F}$, it seems that there should be a physics-based definition of what regions of the boundary are the source, which does not rely on the physics beyond the boundary. Second, by relying on the physics beyond the boundary, we have not gained insight regarding vision, since what is happening beyond the boundary is often not visible from points in free space. We eventually want to understand how the lighting conditions can be inferred by a vision system that is within free space, not outside of it.

With this background and motivation, we now show how to partition a ray manifold into source rays and non-source rays using only the radiance at the boundary of free space. To do so, we introduce a “thought experiment.” Consider a demon that completely absorbs the light arriving at a

given region of $\partial\mathcal{F}$ from outside of \mathcal{F} . Specifically, a given demon absorbs the components of radiance that are emitted and transmitted into \mathcal{F} along a single ray only. We claim that this demon is physically plausible, and we consider the effect of placing demons at the origins of various rays of $\mathcal{M}(\mathcal{F})$.

The basic idea of our definition is this: given a free space \mathcal{F} in a physical scene, define the set of source rays to be the minimal set such that if a demon were placed at the origin of each ray in this set, then the radiance of $\mathcal{M}(\mathcal{F})$ would become identically zero. The relation “minimal” is defined by set inclusion on $\mathcal{M}(\mathcal{F})$.

DEFINITION 2 (Set of Source Rays). Given a free space \mathcal{F} within a physical scene, the set of source rays, $\mathcal{M}_{src} \subseteq \mathcal{M}(\mathcal{F})$, is the minimal set of rays such that if the emitted and transmitted components of each ray in \mathcal{M}_{src} were absorbed at the ray’s point of origin, then the radiance on the manifold $\mathcal{M}(\mathcal{F})$ would become identically zero.

The minimality condition is crucial. For example, if a demon were placed at the origin of every ray in $\mathcal{M}(\mathcal{F})$, then the radiance of $\mathcal{M}(\mathcal{F})$ would surely vanish since no light could enter \mathcal{F} . For most scenes, though, only a subset of $\mathcal{M}(\mathcal{F})$ is needed to make the radiance of $\mathcal{M}(\mathcal{F})$ vanish everywhere. To understand why the minimal subset defines the source, consider an example of a free space consisting of a room with two doorways, one leading outside and the other leading to a closet. Suppose the room also contains a translucent object such as a vase. In order to make the radiance of the room vanish, it is necessary and sufficient to place a demon at each ray that originates in the doorway leading outside. Even though light is transmitted into free space through the vase and through the closet doorway, it is neither necessary nor sufficient to place demons at the origins of these rays to completely darken the room. We argue that this definition of a set of source rays captures our intuition of why, in this example, the closet and vase would not be considered sources.

In the next section, we consider the radiance properties of \mathcal{M}_{src} and show how familiar sources can be characterized by source radiance functions. Before doing so, we make one final point about sources. The set of source rays, \mathcal{M}_{src} , may itself be partitioned into its connected components, and each component considered as a distinct (sub)source. For example, the set of rays originating at a window is a subsource, and the set of rays emitted from a light bulb in the room is another subsource. The two components are topologically disconnected, and it makes sense to think of them as distinct sources. A more interesting example is a room with two windows, as depicted in Fig. 4. If free space is defined to contain the room, as on the left, then the source rays would be those from the sky, which has a single component. If, however, the free space \mathcal{F} were defined to be contained within the room, then the two windows would




FIG. 4. Each connected component of \mathcal{M}_{src} defines a distinct source. In the example of a room with two windows, the connectedness of \mathcal{M}_{src} depends on whether free space contains or is contained within the room.

be two distinct sources. We thus see how the definition of a source depends on how free space is defined. In our view, this dependence is natural.

6. A LIGHT SOURCE HYPERCUBE

When choosing light sources for an environment, it is common to compare sources within a given *type*. For example, lamps are compared by their radiant intensities (W sr^{-1}), which describe how much light they send off to infinity in each direction on the unit sphere. A second example of a type of source is natural daylight. Daylights are compared by the time of day, latitude, weather, etc.

What is rarely done, however, is to compare different *types* of sources with each other. That is, comparing how two lamps differ from each other is clearly separate from comparing how a lamp differs from an overcast sky. What is an appropriate framework for comparing different types of sources? This question is surely relevant for vision, since a visual system often needs to evaluate the lighting conditions in a scene before being able to make sense of the image data. Most computer vision techniques skirt this issue by assuming a certain type of source is present and if necessary estimating its parameters [30]. But before a vision system can infer the parameters of a source, it first needs to identify the type of source(s) present. This problem has simply not been addressed.

In this section, we develop a framework for comparing different types of sources, but ignore the issue of parameterization of a given type of source. (We also do not deal with the question of how to use the framework to identify the type of source in an image.) Our light source framework is based on a class of sources whose parameterizations are trivial, i.e., which have uniform radiance, but whose relationships are nontrivial. Our main observation is that different types of sources can be compared by how light is distributed across the four dimensions of the ray manifold.

The sources we introduce are called *uniform cubic* and are constructed as follows. (It may help the reader to recall Fig. 3.) For a given coordinate plane \mathcal{P}_{z_0} , consider the set of source rays, $\mathbf{r} \in \mathcal{M}_{src}$, that pass through \mathcal{P}_{z_0} and that are restricted to the domain,

$$\varphi_{z_0}^+(\mathbf{r}) = (x, y, p, q) \in \left[-\frac{h_x}{2}, \frac{h_x}{2} \right] \times \left[\frac{h_y}{2}, \frac{h_y}{2} \right] \times \left[-\frac{h_p}{2}, \frac{h_p}{2} \right] \times \left[-\frac{h_q}{2}, \frac{h_q}{2} \right].$$

Each of the parameters h_x , h_y , h_p , and h_q belongs to $(0, \infty)$.

Define the radiance to be uniform over \mathcal{M}_{src} , with value $R(h_x, h_y, h_p, h_q)$. The radiant flux Φ (Watts) of this set of rays, \mathcal{M}_{src} , is

$$\Phi = h_x h_y R(h_x, h_y, h_p, h_q) \int_{-h_p/2}^{h_p/2} \int_{-h_q/2}^{h_q/2} \frac{1}{(1 + p^2 + q^2)^2} dq dp.$$

The factor $(1 + p^2 + q^2)^{-2}$ accounts for the foreshortening of a bundle of rays in direction $(p, q, 1)$.

To compare sources of different dimension, we normalize the radiance to have unit flux, $\Phi = 1$. We do so by defining

$$\alpha(h_p, h_q) = \left(\int_{-h_q/2}^{h_q/2} \int_{-h_p/2}^{h_p/2} \frac{1}{(1 + p^2 + q^2)^2} dp dq \right)^{-1},$$

and letting the uniform radiance be

$$R(h_x, h_y, h_p, h_q) = \frac{1}{h_x h_y} \alpha(h_p, h_q).$$

We also need the windowing function,

$$\omega(u) = \begin{cases} 1, & |u| \leq \frac{1}{2} \\ 0, & |u| > \frac{1}{2} \end{cases}$$

DEFINITION 3. A uniform cubic source of unit flux, centered at position $\mathbf{x} = (0, 0, z_0)$ and direction $(0, 0, 1)$, is a source having radiance function

$$R_{z_0}^+(x, y, p, q) = \frac{1}{h_x h_y} \alpha(h_p, h_q) \omega\left(\frac{x}{h_x}\right) \omega\left(\frac{y}{h_y}\right) \omega\left(\frac{p}{h_p}\right) \omega\left(\frac{q}{h_q}\right).$$

By varying the parameters, h_x , h_y , h_p , h_q , we span a 4-D *light source hypercube*. The corners of this hypercube are shown in Table 1 and are defined by taking the limits of the four parameters to zero or to infinity. Let us discuss a few of these corners.

TABLE 1
The Sixteen Corners of a 4-D Light Source Hypercube

Non-ideal example	Ideal model	h_x	h_y	h_p	h_q	dimension
Overcast sky	Uniform source	∞	∞	∞	∞	4
Cyberware scanner		∞	∞	∞	0	3
Fluorescent tube	Linear source	∞	0	∞	∞	3
Sunlight	Point source at infinity	0	∞	∞	∞	3
	Uniform distribution of rays in a plane	∞	∞	0	0	2
		∞	0	∞	0	2
Louvered linear source (see text)	Fan of rays perpendicular to a linear source	∞	0	0	∞	2
		0	∞	∞	0	2
Small panel light	Point source	0	0	∞	∞	2
Sunlight through crack in doorway	Parallel rays in a plane	∞	0	0	0	1
		0	∞	0	0	1
Rotating spotlight	Fan of rays	0	0	0	∞	1
		0	0	∞	0	1
Spotlight or laser	Single ray	0	0	0	0	0

6.1. Example: A Proximal Point Source

The first corner we discuss is the limit $(h_x, h_y, h_p, h_q) \rightarrow (0, 0, \infty, \infty)$. This limit is the idealization of a small square light source, such as a panel light in a ceiling, as a point light source. We model the limit of h_x and h_y to zero using the Dirac delta function [36], $\delta(t)$, since

$$\lim_{h \rightarrow 0} \frac{1}{h} \omega\left(\frac{u}{h}\right) = \delta(u).$$

To model the limit on h_p and h_q , we observe

$$\lim_{h_p, h_q \rightarrow \infty} \alpha(h_p, h_q) = \pi.$$

Hence, the radiance function of this corner of the hypercube is

$$R_{z_0}^+(x, y, p, q) = \pi \delta(x) \delta(y).$$

We use this model to calculate the irradiance (W m^{-2}) on a planar surface embedded in coordinate plane, \mathcal{P}_{z_1} . Let $\mathbf{x}_1 = (x_1, y_1, z_1) \in \mathcal{P}_{z_1}$, and let $\mathcal{M}_{src}(\mathbf{x}_1)$ denote the set of source rays arriving at \mathbf{x}_1 . We write the irradiance,

$$E(x_1, y_1) = \int_{\mathcal{M}_{src}(\mathbf{x}_1, y_1)} R_{z_1}^+(x_1, y_1, p, q) \frac{1}{(1 + p^2 + q^2)^2} dp dq. \quad (3)$$

Substituting Eq. (1) into Eq. (3) yields

$$E(x_1, y_1) = \int_{\mathcal{M}_{src}(\mathbf{x}_1, y_1)} \delta(x_1 - x_0 - (z_1 - z_0)p) \delta(y_1 - y_0 - (z_1 - z_0)q) \frac{1}{(1 + p^2 + q^2)^2} dp dq,$$

and using the relation, $\delta(at) = (1/a)\delta(t)$, yields

$$E(x_1, y_1) = \frac{(1 + p^2 + q^2)^{-1}}{\|\mathbf{x}_1\|^2},$$

where

$$p = \frac{x_1}{z_1 - z_0} \quad \text{and} \quad q = \frac{y_1}{z_1 - z_0}.$$

This expression for the irradiance produced by a square, panel light is well known [26]. The inverse square law results from the point source approximation, *i.e.*, taking the limit of as h_x, h_y tend to zero. The factor $(1 + p^2 + q^2)^{-1}$ is due to the foreshortening, with respect to the direction $(p, q, 1)$, of the square source and of the surface facet at \mathbf{x}_1 .

6.2. Example: Point Source at Infinity

A second interesting corner of the hypercube is the limit $(h_x, h_y, h_p, h_q) \rightarrow (\infty, \infty, 0, 0)$. This corresponds to a large, collimated set of source rays, which is the model used in classical shape from shading [11] and light source estimation [30]. (Strictly speaking, we do not take h_x, h_y all the way to infinity since this would require that the free space is unbounded. We take instead large bounded values of h_x, h_y .) Observe that

$$\lim_{h_p, h_q \rightarrow 0} \frac{1}{h_p h_q} \alpha(h_p, h_q) \omega\left(\frac{p}{h_p}\right) \omega\left(\frac{q}{h_q}\right) = \delta(p) \delta(q),$$

and so we may write

$$R_{z_0}^+(x, y, p, q) = \frac{1}{h_x h_y} \omega\left(\frac{x}{h_x}\right) \omega\left(\frac{y}{h_y}\right) \delta(p) \delta(q).$$

As in the previous example, we compute the irradiance at $\mathbf{x}_1 \in \mathcal{P}_{z_1}$ by substituting into Eq. (3). This yields

$$E(x_1, y_1) = \frac{1}{h_x h_y} \omega\left(\frac{x}{h_x}\right) \omega\left(\frac{y}{h_y}\right).$$

The positive support of the ω 's corresponds to the ‘‘unshaded’’ region on the plane \mathcal{P}_{z_1} , e.g., the sunbeam on a floor beneath a large window.

Table 1 lists all sixteen (2^4) corners of the hypercube. Most of these corners are intuitive, but a few require clarification. For example, consider the limit $(h_x, h_y, h_p, h_q) \rightarrow (0, 0, 0, \infty)$, which corresponds to a point source that emits or transmits light in a plane only. Such a source would produce a fan of rays. A real example would be a laser beam (or spotlight) that rotates about an axis, e.g., the search lamp of a lighthouse.

A second corner that may be unfamiliar is $(h_x, h_y, h_p, h_q) \rightarrow (\infty, \infty, \infty, 0)$. This is a planar source such that each point in the plane produces a fan of rays perpendicular to the plane. Although no such source exists presently (to our knowledge), a related source does exist and is described as follows. Consider a camera and an ideal linear source (see Table 1) which are mounted on a single moving platform, such that the platform moves in a direction perpendicular to the line defining the linear source. Suppose the camera records only one scanline at a given time and this scanline is parallel to the linear source. The image that would be recorded by the camera would be identical to the image that would be produced by a source having the above radiance function. (This statement ignores surface interreflection effects.) Interestingly, a device similar but not identical to the one just described does exist: the Cyberware laser scanner. This scanner differs from the source just described in that it records images in cylindrical coordinates (θ, y) rather than planar coordinates (x, y) , i.e., the linear source and camera travel around a circle rather than along a line. For an application of this scanner, see [40].

The final corner we consider is the limit $(h_x, h_y, h_p, h_q) \rightarrow (\infty, 0, 0, \infty)$. This is a stationary linear source that emits light only in directions perpendicular to the line of the source. The source is 2-D since for each point on the line there is a circle of ray directions. A source similar to this in common use. Consider a hallway or tunnel that

illuminated by a long fluorescent tube on the ceiling. Surrounding the tube is a sequence of thin disks, called ‘‘louvers,’’ which are like vertebrae surrounding a spinal cord. If the discs were closely spaced and dark, then they would absorb the light from any rays except those that are perpendicular to the source line. The advantage of such a source is that it produces little glare. The source is visible only when looking directly upward.

7. COMPUTATIONAL MODEL

We next turn to the problem of how to compute spatially varying illumination. In this section, we describe data structures and algorithms for performing such a computation. In particular, we consider the problem of recovering shape from shading under spatially varying illumination [18, 19, 20, 38]. Further applications may be found in [16, 17].

We begin with the data structures. Space is discretized as a $N \times N \times N$ cubic lattice. Nodes in this lattice are of three types: SOLID nodes, SURFACE nodes, and FREE nodes. Light is transmitted through FREE nodes. Light is absorbed, reflected, or emitted at SURFACE nodes. Light does not reach SOLID nodes. A labeling of space nodes defines the scene geometry.

For each FREE or SURFACE node, \mathbf{x} , a discretization of the rays at \mathbf{x} is defined by the nodes on a small cube of diameter $2M$ centered at \mathbf{x} . The directions of these rays are determined by nodes on the six faces of this cube. The cube is analogous to the hemicube [6] used in computer graphics. There are two important differences, however. First, the half-width of our cube is much smaller than those used in computer graphics ($M = 5$ vs $M = 50$, typically). This leads to a coarser representation which is less accurate but which is less expensive to compute. Second, our cube is defined at each FREE node, as well as at each SURFACE node, whereas the hemicube is defined only on surfaces. By representing cubes at FREE nodes, we take advantage of the linearity of the coordinate transformations. In particular, we can compute spatially varying illumination using local operations, namely coordinate chart transformations, instead of the global visibility operations used in computer graphics.

Coordinate charts on this discretized ray manifold are illustrated Fig. 5 (left). For a given face F of the ray cube (i.e., there are 6 faces), consider the i th plane in the cubic space lattice parallel to F . Let \mathcal{U}_i^F denote the set of rays whose directions are contained in F and who pass through FREE or SURFACE nodes in plane i . As in the continuous case, we define the coordinate chart φ_i^F to be a mapping from \mathcal{U}_i^F into \mathfrak{R}^4 such that

$$\mathcal{U}_i^F \equiv \{(\varphi_i^F)^{-1}(x, y, p, q) : (x, y, i) \text{ is a FREE or SURFACE node, and } (p, q, M) \in F\}.$$




FIG. 5. (Left) A coordinate chart on the discretized ray manifold. (Right) A local coordinate transformation.

Neighboring coordinate charts typically overlap (see Fig. 5, right). For example, consider a FREE node $\mathbf{x} = (x, y, i)$, and a ray passing through this node in direction $\mathbf{u} = (p, q, M)$. This ray could be parameterized by either

$$(\varphi_i^F)^{-1}(x, y, p, q) \in \mathcal{U}_i^F \quad \text{or}$$

$$(\varphi_{i+M}^F)^{-1}(x + p, y + q, p, q) \in \mathcal{U}_{i+M}^F.$$

A basic operation in computing spatially varying illumination is to transform the radiance in one coordinate chart to that of a neighboring coordinate chart. For example, this allows us to compute the radiance in coordinate plane $i + M$ given the radiance in plane i . Let $R_i^F(x, y, p, q)$ denote the radiance of rays in chart φ_i^F . Then, given R_i^F , we may compute R_{i+M}^F as follows:

Coordinate_Transformation (i, F)

```
{ for all  $(x, y)$ 
   $\mathbf{x} := (x, y, i)$ ;
  for all  $\mathbf{u} = (p, q, M) \in F$ 
    if  $(\mathbf{x}, \mathbf{x} + \mathbf{u}) \subset \mathcal{F}$ 
       $R_{i+M}^F(x + p, y + p, p, q) := R_i^F(x, y, p, q)$ ;
    else Local_Tunnel( $\mathbf{x}, \mathbf{x} + \mathbf{u}$ );
}
```

The subroutine **Local_Tunnel**($\mathbf{x}, \mathbf{x} + \mathbf{u}$) is needed for situations in which the line segment $(\mathbf{x}, \mathbf{x} + \mathbf{u})$ passes through a SURFACE node. When this occurs, the radiance from the nearest such node is used to determine the radiance $R_{i+M}^F(x + p, y + p, p, q)$.

Since the radiance at plane $i + M$ is computable from the radiance at plane i , we can compute the radiance of any plane provided we know the radiance in the first M planes. Then, once we compute the radiance of a given plane, $i + M$, from plane i , we may discard the radiance of plane i . This allows us to keep track only of M consecutive

planes at a time, and hence allows us to reduce the memory required to compute an entire sequence of coordinate transformations. (For a detailed analysis of the space and time costs of the algorithm, see [16].)

7.1. Example: Shape-from-Shading

We now apply the above algorithm for computing coordinate transformations to the shape-from-shading problem. Recall that in classical shape-from-shading, it is assumed that there is a single point source at infinity [11]. As we have shown throughout this paper, in many situations, light sources produce an illumination that varies across space. This raises the question of how shape may be computed from shading under spatially varying illumination.

In [18, 19], we introduced an algorithm that solves this problem in the case of a uniform hemispheric source, which approximates the sky on a cloudy day. The algorithm is based on the space lattice data structure and coordinate transformations discussed above. (More recently, an alternative data structure and algorithm was used to solve the problem [38].) The key idea was that, under uniform diffuse lighting, the direct irradiance at a point is determined primarily by the solid angle of the source that is visible from that point. This solid angle varies along a surface. For example, concavities tend to be dark, while hilltops tend to be bright. It was shown that the solid angle effect, which is essentially shadowing, typically dominates surface normal effects.

In [20], we generalized that algorithm to address the case of proximal sources, such as a lamp or spotlight. Such sources are very common both in indoor scenes and in outdoor scenes at night. For example, consider a dark roadway illuminated by car headlights or by a street lamp. Vision systems need to be able to use the shading pattern on the road to infer surface shape, in particular, the slant of a wall or the curve of a roadway.

As a first step toward solving this new shape-from-shading problem, we consider a simplified version of it in which the scene is a slanted plane, illuminated by a small spherical light source at the viewer, and viewed orthographically. The algorithm we use to recover the plane is an extension of the original “cloudy day” algorithm [18, 19]. First, nodes at depth $z = 0$ are considered. The lighting conditions are assumed to be known, and hence the radiance in coordinate planes, $z = \{0, \dots, M - 1\}$, is known. The algorithm proceeds by induction. Given the radiance at each node at depth, k , decide which of these nodes are SURFACE nodes by comparing the measured image intensity, $I(x, y)$, with a model, $I_{\text{model}}(x, y)$. For the case of a proximal source, our intensity model is

$$I_{\text{model}}(\mathbf{x}) = \sum_{\mathbf{u} \in \mathcal{U}_{\text{src}}(\mathbf{x})} R(\mathbf{x}, \mathbf{u}) \Delta\Omega,$$

where $\Delta\Omega$ is the angle subtended by a particular node on the hemicube. (Note that this model ignores surfaces normal and interreflection effects and serves merely to illustrate the basic idea of the approach.) Then, for each FREE node at depth $(k + 1)$, compute the radiance at that node. Iteration on k eventually results in each pixel (x, y) having a computed SURFACE node at a depth $z(x, y)$.

Shape from Shading $(I(x, y))$:

```
{
  for all  $(x, y)$ ,
     $z(x, y) := \infty$ ;
  for all  $\mathbf{u} \in F$ ,
    initialize source radiance,  $R_0(x, y, \mathbf{u})$ ;
   $i := 0$ ;
  repeat
    for all  $(x, y)$ ,
       $\mathbf{x} := (x, y, i)$ ;
      if  $z(x, y) = \infty$  and  $|I_{\text{model}}(x, y) - I(x, y)| < \epsilon$ 
        then  $z(x, y) := i$ ;
      Coordinate_Transformation  $(i, F)$ 
       $i := i + 1$ ;
  until  $z(x, y) < i$  for all  $(x, y)$ 
}
```

Figure 6 shows the rendered plane along with the computed depth maps. The upper row shows three rendered images, obtained from different source radiance functions. These correspond to a light bulb (left), a weakly directed source (middle), and a spotlight (right). Observe that as the source radiance becomes more directed toward the optical axis of the camera, the image intensity maximum shifts toward the center of the image. The quantization artifacts in the rendered images are the result of the coarse discretization of the cube of ray directions. (The reader should blur his/her eyes to properly view the images.)

A depth map was then recovered from each of the three images using the above algorithm. For each case, a surface approximating a slanted plane was recovered. The corners of the rightmost example illustrate an important ambiguity in the algorithm. Since a spotlight gives off a cone of light, the columns of free space at the image corners are darkest at shallow points (above the source cone), brighter as the depth increases, and darker again as the distance from the source increases. The algorithm cannot distinguish the two causes of darkness.

The example illustrates our basic approach to computing shape from shading under spatially varying illumination. More robust data structures and algorithms are needed, however. For example, an alternative data structure and algorithm for solving the cloudy day problem was recently presented [38] and was shown to recover surface shape more accurately than the original algorithm of [18, 19]. The new algorithm could be extended to the case of a proximal source as well. Because the new algorithm accounts for surface normal variations as well as spatially varying illumination, we would expect it to perform better than the algorithm above.

Of course, there are a number of issues which the algorithm above does not address; for example, What if the surfaces are not Lambertian? How should interreflection effects be modeled? How can the lighting conditions be inferred if they are not yet known? To build robust computer vision systems, these questions also need to be answered.

8. THE ECOLOGICAL IMPLICATIONS OF ILLUMINATION

It is a common observation that cloudy days are “gray days”; that is to say, shadows are less distinct than on sunny days, and contrast is, on average, lower. This observation is related to the light-source hypercube just developed, in that “gray days” are days in which the source is 4-D, while sunny days enjoy sources that are lower dimensional. This relationship between the dimensionality of sources and the character of light provides a language for discussing complicated lighting in natural situations. To illustrate, imagine a laser beam shining on a wall in a room. In this example a 0-D source scatters on the wall to form a 2-D source, which, through further mutual reflections, would illuminate the rest of the room as though it were a 4-D source. An observer viewing the room without seeing either the laser or its first reflection would thus see a room much as it would appear from a diffuse source. (Of course, much artificial lighting is designed to produce this diffusing effect, as when a lamp shade turns a filament into a 4-D source.) Should the observer’s viewpoint shift in the above example, to bring the laser spot on the wall into view, then this spot would appear so much brighter than the rest of




FIG. 6. (Upper row) Three synthetic images of a slanted plane illuminated by a spherical source. On the left, the radiance is isotropic, e.g., a bare light bulb. On the right, the radiance is strongly directed toward the viewing direction. (Lower row) The recovered depth maps (see text).

the wall, and the viewer would infer it to be the source [41, 5]. Another scenario: consider a thin beam of sunlight entering a room through a doorway. This is a 1-D source that would become 3-D after reflecting off a wall. Again, from a viewer's perspective, whether the source is 1-D, 3-D, or 4-D depends on whether the viewer sees the sun through the doorway or the line of sunlight reflected from the wall. One final example: it is often observed that the light distribution in a rain forest is complicated. But by defining free space within the forest, our formal definition implies that, since light is arriving from basically all directions and all angles, it is essentially the same as the 4-D situation. Through these examples, we see how the complexity of lighting is clarified by understanding the geometry and dimension of sources.

From an ecological perspective there are many reasons why illumination information can be useful. First, back lighting is common in natural scenes, for example, when an object is viewed against an overcast sky or against a window, or toward a sunrise or sunset. Back lighting produces high image contrast. As a result, surface properties tend to be poorly visible, as is illustrated with the Rubin's vase cartoon of Fig. 7. When the scene appears to consist of two faces in profile, the faces appear to be back lit. The albedos of the faces are uncertain since they lie in the attached shadow of an extended source. Note that a vision system should not conclude that the albedo of the faces is low relative to the background simply because the image

intensities are lower, but rather should reason that the albedo is uncertain because the surfaces are back lit. More directly to the point of this paper, notice how, when a figure-ground reversal occurs in which the vase becomes the figure, there is also a lighting reversal. The vase now appears to be illuminated from the foreground, e.g., by a spotlight, and the background appears dark because it lies in shadow.

Another example is that shadows must be distinguished from surfaces having low reflectance. This is important for




FIG. 7. When the image is perceived as two faces in profile, the albedo of the faces is uncertain since the scene is back lit. When there is a figure-ground reversal, so that a central vase is perceived as figure, there is also a lighting reversal. The vase appears to be illuminated from the foreground, whereas the background is dark because it lies in shadow.

object recognition; *e.g.*, a moving cast shadow of a swaying branch should be distinguished from a moving dark-colored animal. It is also important because shadows, as dark 3-D regions, are good hiding places. If an animal were to mistake a dark-colored but well-illuminated patch on the ground for a shadow and move into that patch, then it would increase its visibility rather than decrease it! Similarly, an animal should distinguish shadows from dark-colored regions to avoid stepping in holes in the ground and to locate cool regions of the scene to rest, *e.g.*, for regulating body temperature.

Our next examples relate to the observation that different illumination conditions reveal different properties of shape. We previously observed that, under diffuse illumination, “darker does not mean deeper” [18, 19]. For another example, consider a matte surface terrain with a low angle of illumination, *e.g.*, sunrise. The image intensity variations will be biased toward the direction of the source [31]. In particular, a surface region in which the hills and valleys are elongated and parallel to the source would produce little shading. A vision system that is able to estimate the illuminant direction [30] should also represent the resulting uncertainty in shape. A similar lesson holds for other lighting conditions. For example, under diffuse lighting or even under overhead lighting, shallow variations in surface depth produce little shading. Illumination representations are a basis for reasoning about which cues are visible and which are not.

Our final comment about the model is not in the form of an example, but is rather an attempt to approach the biology of illumination representations, should they exist. Recall that in the Introduction we used the relationship between illumination and space as motivation, observing that, *e.g.*, lighting in restaurants creates a very different atmosphere from lighting in stadiums. The above examples provide further evidence that illumination is important, which raises the difficult question of how animals might represent such information. It is, of course, entirely premature to make any specific statements on this point here, since no direct experiments have been performed to our knowledge. However, we are intrigued by one final observation which, since it is motivated by our model, could structure first thoughts on the problem.

The heart of our model is that illumination is a representation based on rays in free space, and it is this relationship that relates illumination to primate physiology. Without going into details (see references below), there are different “streams” for the processing of visual information. Both begin in the occipital cortex, and one proceeds to the temporal cortex while the other proceeds to the parietal cortex. Mishkin and Ungerleider suggest that these represent “what” and “where” pathways [42], while Goodale and Milner suggest they represent “perception” vs “action” [25]. The anatomy involved in this distinction is clear,

although the interpretation still remains controversial. However, common through both of these interpretations is that the parietal stream involves information about space, while the temporal stream involves more traditional visual representations. For robotics, the information about visuomotor control would be characterized to support path and trajectory planning over free spaces. We are struck by the observation that our model for illumination is also defined over (free) space and is thus appropriately placed in the parietal stream. Perhaps this connection to movement and visuomotor control is the basis for connecting lighting to perceived space. Much clearly remains to be done before such speculations could be refined to testable hypotheses, but the character of the match remains suggestive at this time. The speculation that illumination is a parietal property would open a totally new arena for physiological experimentation.

9. CONCLUSION

Vision is the process by which the brain interprets light. Light is central to vision, as sound is to hearing, and pressure is to touch. Understanding how a vision system represents and computes illumination is therefore central to understanding how vision works. Unfortunately, illumination has been given a secondary role in vision research. We believe there are two reasons for this: (1) computational constraints on spatially varying illumination have been poorly understood; without such constraints, a computational theory was impossible; (2) the ecological importance of recovering illumination has been downplayed relative to that of recovering shape, material, and 3-D layout. Our contribution in this paper is therefore the following. First, we have provided a computational model of spatially varying illumination, *i.e.*, what it is, and constraints on how to represent and compute it. Admittedly, our model is incomplete, *e.g.*, we do not yet understand how the lighting conditions of a scene may be inferred from images. However, the model provides a solid foundation on which more comprehensive and higher-level models can be built. With this in mind, we have also discussed ecological reasons for why illumination should be represented, namely, representing the illumination allows a vision system to reason about which parts of the scene are visible and why. In closing, we observe that illumination has enjoyed a special position in architectural design, but has been avoided in physiology and grossly oversimplified in computer vision. We hope our model will cast sufficient light on illumination to bring it into the mainstream of vision research.

APPENDIX

Let $\mathcal{U}_{z_0}^- \subset \mathcal{M}(\mathcal{F})$ denote the set of rays in $\mathcal{M}(\mathcal{F})$ that pierce the coordinate plane \mathcal{P}_{z_0} in the negative z direction. If $\mathbf{r} \in \mathcal{U}_{z_0}^-$ is the ray that passes through \mathbf{x}_0 in direction $(p, q, -1)$, then define

$$\varphi_{z_0}^-(\mathbf{r}) \equiv (x_0, y_0, p, q).$$

The coordinate charts $\varphi_{x_0}^+$, $\varphi_{x_0}^-$, $\varphi_{y_0}^+$, $\varphi_{y_0}^-$ are defined similarly.

DEFINITION 4. For a given $\mathcal{M}(\mathcal{F})$, let

$$\mathcal{A} \equiv \bigcup_{\mathbf{x}_0 \in \mathcal{F}} \{\varphi_{x_0}^+, \varphi_{x_0}^-, \varphi_{y_0}^+, \varphi_{y_0}^-, \varphi_{z_0}^+, \varphi_{z_0}^-\}$$

be an atlas of coordinate charts defined by all points in free space.

PROPOSITION 1 (Ray manifold). *Given a free space \mathcal{F} , the pair $(\mathcal{M}(\mathcal{F}), \mathcal{A})$ is a 4-D manifold.*

Proof. The domains of the charts in \mathcal{A} cover $\mathcal{M}(\mathcal{F})$ since any ray $\mathbf{r} \in \mathcal{M}(\mathcal{F})$ passes through at least one point in free space, and hence pierces at least one of the six coordinate planes at that point. Moreover, the domains of the charts are open sets. To see this, consider the ray \mathbf{r} in Fig. 3. Since $\mathbf{x}_0 \in \mathcal{F}$, there is an ε -ball in \mathbb{R}^3 that is centered at \mathbf{x}_0 and contained in \mathcal{F} . The rays that pass through this ε -ball and whose directions have a positive z component define an open neighborhood of \mathbf{r} .

All that remains to be shown is that the intersection of coordinate charts is well-behaved. We have already considered the case of $\varphi_{z_0}^+$, $\varphi_{z_1}^+$, so only two further cases need to be considered:

- $\varphi_{z_0}^+$, $\varphi_{z_1}^-$: The domains of these charts do not intersect.
- $\varphi_{z_0}^+$, $\varphi_{x_1}^+$: Let $\mathbf{r} \in \mathcal{U}_{z_0}^+ \cap \mathcal{U}_{x_1}^+$, and again let $\varphi_{z_0}^+(\mathbf{r}) \equiv (x_0, y_0, p, q)$. Observe that $p \neq 0$, since otherwise \mathbf{r} would be strictly contained in the plane \mathcal{P}_{x_1} rather than transverse to it and hence would not belong to $\mathcal{U}_{x_1}^+$. If $(x_0, y_0, z_0) \neq (x_1, y_1, z_1)$ then

$$(p, q, 1) = \left(\frac{x_1 - x_0}{z_1 - z_0}, \frac{y_1 - y_0}{z_1 - z_0}, 1 \right).$$

Moreover,

$$\varphi_{x_1}^+(\mathbf{r}) = \left(y_1, z_1, \frac{y_1 - y_0}{x_1 - x_0}, \frac{z_1 - z_0}{x_1 - x_0} \right) = \left(y_1, z_1, \frac{q}{p}, \frac{1}{p} \right),$$

so that

$$\varphi_{x_1}^+ \circ (\varphi_{z_0}^+)^{-1}(x_0, y_0, p, q) = \left(y_0 + \frac{q}{p}(x_1 - x_0), z_0 + \frac{x_1 - x_0}{p}, \frac{q}{p}, \frac{1}{p} \right).$$

The same formula applies in the case that $(x_0, y_0, z_0) = (x_1, y_1, z_1)$. Since the domains of the charts are open sets and since the coordinate transformations are continuous, the proposition follows.

Note that the proposition holds even if a free space is non-convex. To appreciate this, consider two simple examples of sets that fail to be manifolds, namely the letters ‘‘T’’ and ‘‘X.’’ Each is a union of 1-D manifolds (curves). However, neither letter is itself a 1-D manifold. The failure occurs because there is no neighborhood of the junction point that is homeomorphic to \mathbb{R} . Specifically, removing a point from \mathbb{R} disconnects it into two components, whereas removing the junction point from either letter disconnects the letter into more than two components. The example is relevant to $\mathcal{M}(\mathcal{F})$ since T-junctions are well known to occur in images when free spaces are non-convex.

ACKNOWLEDGMENTS

This research was supported by grants from NSERC and AFOSR, while the authors were at the Center for Intelligent Machines, McGill

University, Montreal, Canada. The authors thank Steve Shafer, Allan Jepson, James Elder, Peter Belhumeur, David Kriegman, Kaleem Siddiqui, and Lance Williams for their helpful comments.

REFERENCES

1. E. H. Adelson and J. R. Bergen, The plenoptic function and the elements of early vision, in *Computational Models of Visual Processing*, MIT Press, Cambridge, MA, 1990.
2. J. Arvo and D. Kirk, Fast ray tracing by ray classification, *Comput. Graphics* **21** (1987), 55–64.
3. P. N. Belhumeur and D. J. Kriegman, What is the set of images of an object under all possible lighting conditions? in *IEEE Conference on Computer Vision and Pattern Recognition, San Francisco, June 1996*.
4. A. Blake and G. Brelstaff, Geometry from specularities, in *Proceedings 2nd ICCV, Tarpon Springs, FL, 1988*, pp. 394–403.
5. F. Bonato and A. L. Gilchrist, The perception of luminosity on different backgrounds and in different illuminations, *Perception* **23** (1994), 991–1006.
6. M. F. Cohen and D. P. Greenberg, The hemicube: A radiosity solution for complex environments, *Comput. Graphics* **19** (1985), 31–40.
7. J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes, *Computer Graphics: Principles and Practice*, 2nd ed., Addison–Wesley, Reading, MA, 1990.
8. D. Forsyth and A. Zisserman, Reflections on shading, *IEEE Trans. Pattern Anal. Machine Intelligence* **13** (1991), 671–679.
9. J. Gibson, *The Senses Considered as Perceptual Systems*, Houghton Mifflin, Boston, 1966.
10. B. K. P. Horn, Understanding image intensities, *Artificial Intelligence* **8** (1977), 201–231.
11. B. K. P. Horn and M. J. Brooks, Eds., *Shape from Shading*, MIT Press, Cambridge, MA, 1989.
12. B. K. P. Horn and R. W. Sjoberg, Calculating the reflectance map, *Appl. Opt.* **18** (1979), 1770–1779.
13. K. Ikeuchi, Determining surface orientations of specular surfaces by using the photometric stereo method, *IEEE Trans. Pattern Anal. Machine Intelligence* **3** (1981), 661–669.
14. J. J. Koenderink and A. J. van Doorn, Photometric invariants related to solid shape. *Opt. Acta.* **27** (1980), 981–996.
15. E. H. Land and J. J. McCann, Lightness theory, *J. Opt. Soc. Am.* **61** (1971), 1–11.
16. M. S. Langer, P. Breton, and S. W. Zucker, Massively parallel radiosity in the presence of multiple isotropic volume scattering, in *Graphics Interface '95, Quebec City, Canada, May 1995*, pp. 103–108.
17. M. S. Langer, G. Dudek, and S. W. Zucker, Space occupancy using multiple shadowimages, in *IROS '95, International Conference on Intelligent Robotics and Systems, Pittsburgh, Aug. 1995*.
18. M. S. Langer and S. W. Zucker, Diffuse shading, visibility fields, and the geometry of ambient light, in *Proceedings of the 4th International Conference on Computer Vision, Berlin, Germany, 1993*, pp. 138–147.
19. M. S. Langer and S. W. Zucker, Shape-from-shading on a cloudy day, *J. Opt. Soc. Am. A* **11** (1994), 467–478.
20. M. S. Langer and S. W. Zucker, Spatially varying illumination: A computational model of converging and diverging sources, in *Computer Vision—ECCV'94*, (J.-O. Eklundh, Ed., Vol. 2, pp. 227–232, Springer-Verlag, Berlin/New York, 1994.
21. M. Levoy and P. Hanrahan, Light field rendering. in *Proceedings of SIGGRAPH, Aug. 1996*.

22. A. K. Mackworth, Interpreting pictures of polyhedral scenes, *Artificial Intelligence* **4** (1973), 121–137.
23. D. Marr, *Vision*, pp. 24–29, Freeman, San Francisco, 1982.
24. L. McMillan and G. Bishop, Plenoptic modelling: An image-based rendering system, in *SIGGRAPH Proceedings, 1995*, pp. 39–46.
25. A. D. Milner and M. A. Goodale, *The Visual Brain in Action*, Oxford Univ. Press, London, 1995.
26. P. H. Moon, *The Scientific Basis of Illuminating Engineering*, Dover, New York, 1961.
27. S. K. Nayar and R. M. Bolle, Reflectance based object recognition, *Int. J. Comput. Vision* **17** (1996), 219–240.
28. S. K. Nayar, M. Watanabe, and M. Noguchi, Real-time focus range sensor, in *Proceedings of the International Conference on Computer Vision, Cambridge, MA, June 1995*, pp. 995–1001.
29. J. Oliensis and P. Dupuis, A global algorithm for shape from shading, in *Proceedings of the 4th International Conference on Computer Vision, Berlin, Germany, 1993*, pp. 692–701.
30. A. Pentland, Finding the illuminant direction, *J. Opt. Soc. Am.* **72** (1982), 448–455.
31. A. Pentland, Linear shape from shading, *Int. J. Comput. Vision* **4** (1990), 153–162.
32. A. Pentland, S. Scherrock, T. Darrell, and B. Girod, Simple range cameras based on focal error, *J. Opt. Soc. Am. A* **11** (1994), 2925–2934.
33. M. Planck, *The Theory of Heat Radiation*, Dover, New York, 1959.
34. I. M. Singer and J. A. Thorpe, *Lecture Notes on Elementary Topology and Geometry*, Springer-Verlag, Berlin/New York, 1967.
35. D. M. Y. Sommerville, *Analytic Geometry of Three Dimensions*, Cambridge Univ. Press, Cambridge, UK, 1959.
36. H. Stark and F. B. Tuteur, *Modern Electrical Communications*, Prentice Hall, Englewood Cliffs, NJ, 1979.
37. G. R. Steffy, *Architectural Lighting Design*, Van Nostrand–Reinhold, New York, 1990.
38. A. J. Stewart and M. S. Langer, Towards accurate recovery of shape from shading under diffuse lighting, in *IEEE Conference on Computer Vision and Pattern Recognition, San Francisco, June 1996*.
39. S. J. Teller, *Visibility Computations in Densely Occluded Polyhedral Environments*, Ph.D. thesis, University of California at Berkeley, 1992.
40. N. F. Troje and H. H. Bülthoff, Face recognition under varying poses: The role of texture and shape, *Vision Research* **36** (1996), 1761–1771.
41. S. Ullman, On visual detection of light sources, *Biol. Cybernetics* **21** (1976), 205–212.
42. L. G. Ungerleider and M. Mishkin, Cortical systems, in *Analysis of Visual Behavior*, pp. 549–586, MIT Press, Cambridge, MA, 1982.