What is a light source?

Michael S. Langer
NEC Research Institute
4 Independence Way
Princeton, NJ 18540

Steven W. Zucker
Depts. of Comp. Sci. and Elec. Eng.
Yale University
New Haven, CT 06520-8285

Abstract

Traditional light source modelling is concerned with specific types of light sources, the two most common of which are point sources and daylight. Little attempt has been made, however, to relate different types of sources to each other. For example, how may the lighting from an overcast sky be compared to that from a lamp? Having a theoretical framework to compare different types of light sources is important for computer vision, in particular for understanding shading and shadow cues. A vision system needs to take account of the light source in order to interpret these cues. In this paper, we present a framework for comparing types of light sources which is based on a dimensional analysis of the set of light rays in a free space. Specifically, we introduce a 4-D light source hypercube in which the different types of sources may be embedded and compared. We also present a novel definition for light sources which generalizes the standard definition of a source as an emitter.

1 Introduction

Many types of lighting models have been used in computer vision. Examples include a point source at infinity[2], a proximal source, a linear source [4], a uniform hemispheric source [5], a sum of point sources at infinity [1]. Unfortunately, there presently exists no framework to relate these models. This is problematic since many vision algorithms – in particular, shape from shading – are designed for a specific type of light source, such as a point source at infinity [3] or a uniform hemispheric source[5]. Determining which algorithm to apply involves determining which type of source is present.

Our goal in this paper is to take a first step towards solving this general light source estimation problem, by providing a theoretical framework in which different types of sources can be compared. Our goal differs from traditional goals of light source modelling in that we compare different types of light sources to each other, rather than comparing different instances of a single type. To appreciate this distinction, consider two common types of source: the proximal point source, and daylight. Different proximal point sources are compared in terms of their radiant intensity (W sr^{-1}), which is how much power each source sends to infinity in each direction. (For example, the radiant intensity of a lamp may depend on the filament, the lamp shade, the reflector, etc.) The definition of radiant intensity [8] does not apply to daylight, however. Instead, daylights are compared only to each other *e.g.* by the time of day, the latitude, the weather, etc. [8]. How can the lighting from a proximal point source be compared to that from daylight?

2 Overview

Our idea for comparing sources is to perform a dimensional analysis on the set of light rays in a scene. This set of rays, which we call the *ray manifold*, was introduced in [6]. In Section 3 we review the construction of the ray manifold, and in Section 4 we use the manifold to carry out a dimensional analysis of sources. In particular, we derive a 4-D *light source hypercube*, in which each corner is a source type. These sixteen (2^4) corners capture most natural sources, as well as many artificial ones.

In Section 5, we specify rigorously what we mean by the term "light source," by providing a definition in terms of a partition of the ray manifold. This definition generalizes the standard definition of a source as an emitter, and also clarifies a number of interesting puzzles about sources.

3 The Ray Manifold

Our framework for modelling light sources is based on the set of light rays in a scene. The reason for using rays is a well-known empirical law of radiometry that, in the absence of atmospheric emission or scattering and above the scale of diffraction, radiance is constant along a geometric ray [10]. The law is central to the physics of light, and is the main reason that rays are used in geometric optics.

To formally specify the set of rays in a scene, we begin by defining a **free space**, $\mathcal{F}$, to be a bounded, open, connected subset of $\Re^3$, which has a piecewise smooth boundary, $\partial\mathcal{F}$, and which is free of objects. By an "object", we mean anything that scatters, reflects, emits or absorbs light. For example, the interior of a candle flame or a cloud cannot be part of free space by this definition, since these objects emit and scatter light, respectively.

A free space may be non–convex. In particular, a free space may have holes (objects) in which light is reflected, emitted, absorbed or scattered. Moreover, the boundary of a free space, $\partial\mathcal{F}$, need not be restricted to the surfaces of objects. For example, a window or doorway could be part of $\partial\mathcal{F}$, as could an imaginary bubble enclosing a free space in an outdoor scene.

We now define the set of rays, $\mathcal{M}(\mathcal{F})$, in a given free space. For any two points $\mathbf{x}_1, \mathbf{x}_2 \in \Re^3$, 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 across free space, 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 defined 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.

We parameterize the set of rays $\mathcal{M}(\mathcal{F})$ as follows. Consider a plane

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

Let $\mathbf{r}$ be a ray that pierces this plane in the positive z direction, such that $\mathbf{r}$ passes through $\mathbf{x}_0 = (x_0, y_0, z_0) \in \mathcal{F}$, and in direction $(p_0, q_0, 1)$. We may parameterize $\mathbf{r}$ using four coordinates, (x_0, y_0, p_0, q_0).

This defines a local coordinate system on the set of rays $\mathcal{M}(\mathcal{F})$. The coordinate system is "local" because not all rays intersect $\mathcal{P}_{z_0}$. By defining a family of such local coordinate systems using different planes, in particular planes parallel to the y and z directions, one may prove the following [6].

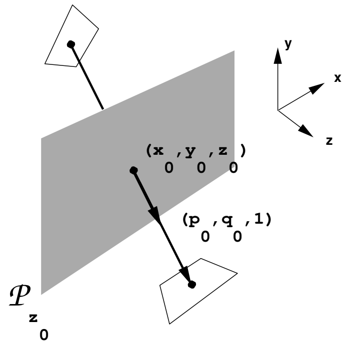

Figure 1: A ray $\mathbf{r}$ passing through (x_0, y_0, z_0) and in direction $(p_0, q_0, 1)$. This ray may be parameterized as (x_0, y_0, p_0, q_0).

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

A key observation is that, because radiance R is constant along a ray in free space,

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

Specifically, let $R^+_{z_0}(x_0, y_0, p_0, q_0)$ denote the radiance of the ray $\mathbf{r}$, and suppose that $\mathbf{r}$ also passes through a plane $\mathcal{P}_{z_1}$. (The $^+$ superscript of $R^+_{z_0}$ means that the ray is in direction $(p_0, q_0, 1)$ rather than in direction $(p_0, q_0, -1)$.) From Proposition 1 above, we have

$$\begin{aligned} R^+_{z_0}(x_0, y_0, p_0, q_0) \quad &= \\ R^+_{z_1}(x_0 + (z_1 - z_0)p_0, \ y_0 + (z_1 - z_0)q_0, \ p_0, q_0) \ . \quad &(1) \end{aligned}$$

We will use this relation extensively the following section, when we examine the irradiance produced by different types of sources.

4 The Light Source Hypercube

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. In Section 5, we will define such a partition. For now, let us assume this can be done, and let $\mathcal{M}_{src}$ denote the set of source rays, where $\mathcal{M}_{src} \subseteq \mathcal{M}(\mathcal{F})$.

Our goal is to develop a framework for comparing different types of sources. For simplicity, we ignore the traditional concerns of light source modelling, namely, how to parameterize sources of a given type. Instead, we introduce a class of sources which span many types, but whose parameterizations are trivial: *each source has uniform radiance.* What is non-trivial about these sources is how light is distributed over the four dimensions of the ray manifold. The key point is that, by restricting the distributions to lie in sets of various dimensions, we define different types of sources.

4.1 Uniform Cubic Sources

The class of sources is constructed as follows. (Recall Figure 1). Define four parameters, h_x, h_y, h_p, and h_q, each of which belong to $\Re^+$. For a given plane $\mathcal{P}_{z_0}$ (recall Fig. 1), define a set of source rays,

$$\mathcal{M}_{src} \equiv \{ (x,y,p,q) : x \in [-\tfrac{h_x}{2}, \tfrac{h_x}{2}], \quad y \in [-\tfrac{h_y}{2}, \tfrac{h_y}{2}], p \in [-\tfrac{h_p}{2}, \tfrac{h_p}{2}], q \in [-\tfrac{h_q}{2}, \tfrac{h_q}{2}] \}.$$

This set is a 4-D hypercube.

For simplicity, let the radiance of M_{src} be uniform with value $R(h_x,h_y,h_p,h_q)$ to be defined below. It is straightforward to show that the radiant flux (Watts) of the rays in M_{src} is

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

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

To compare different types of sources, we normalize $R(h_x,h_y,h_p,h_q)$ so that each source has unit flux. We do so by defining

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

and defining the uniform radiance to be

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

Finally, we define the windowing function,

$$\omega(u) \equiv \begin{cases} 1, & |u| \le \frac{1}{2} \\ 0, & |u| > \frac{1}{2} \end{cases} .$$

Definition 2 *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,*

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

Observation 1 *The four parameters,* h_x*,* h_y*,* h_p*,* h_q *define a 4-D light source hypercube.*

The corners of this hypercube are shown in Table 1. These corners are defined by taking the limits of the four parameters to zero or to infinity. We now discuss these corners in detail.

4.2 Example 1: a proximal point source

The first corner we discuss is the limit $(h_x,h_y,h_p,h_q) \to (0,0,\infty,\infty)$. This limit is the point source idealization of a small square light source such as a panel light in a ceiling. We model the limit of h_x and h_y to zero using the Dirac delta function [11], $\delta(t)$, by observing

$$\lim_{h \to 0} \ \frac{1}{h} \, \omega(\frac{u}{h}) = \delta(u) \ .$$

To model the limit on h_p and h_q, we observe

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

Hence, the radiance is

$$R^+_{z_0}(x,y,p,q) = \pi \ \delta(x) \ \delta(y) \ . \tag{2}$$

As a sanity check, let us calculate the irradiance (W m^{-2}) produced by this source onto a planar surface, $\mathcal{P}_{z_1}$. 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, $E(x_1,y_1)$, as

$$\int_{\mathcal{M}_{src}(x_1,y_1)} R^+_{z_1}(x_1,y_1,p,q) \frac{dp \ dq}{(1+p^2+q^2)^2} \ . \tag{3}$$

Substituting Equations (1) and (2) into (3) yields

$$E(x_1,y_1) = \int \delta(x_1 - x_0 - (z_1 - z_0) \ p) \ \delta(y_1 - y_0 - (z_1 - z_0) \ q \ ) \frac{dp \ dq}{(1+p^2+q^2)^2}.$$

Using the relation, $\delta(at) = \frac{1}{a} \ \delta(t)$, yields

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

where $p = \dfrac{x_1}{z_1 - z_0}$ and $q = \dfrac{y_1}{z_1 - z_0}$.

This expression for the irradiance produced by a square panel light is well-known [8]. The inverse square law results from the point source approximation *i.e.* taking the limit 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 $\mathbf{x}_1$.

4.3 Example 2: a 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 [3] and light source estimation [9]. (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

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

and so we may write,

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

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

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

The positive support of the ω's corresponds to the UN-shadowed region on the plane $\mathcal{P}_{z_1}$, *e.g.* the sunbeam on a floor beneath a large window.

4.4 Further Examples

Table 1 lists all sixteen (2^4) corners of the hypercube. Most of these are intuitive, but a few require clarification (see also [7]). 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.* a search lamp.

A second unfamiliar corner 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 is in common use. Consider a hallway or tunnel 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 upwards.

One final point about dimension. Strictly speaking, physical sources are always 4-D, never smaller. Taking the limit of $h \rightarrow 0$ is a mathematical abstraction. This is in exactly the same spirit as when one refers to a line drawn on a page as a 1-D curve. (A line drawn with a physical pen always has a finite thickness, and hence is in reality a 2-D set.) By treating sources in this idealized way, *i.e.* as a set of dimension less than or equal to four, we can articulate the essential differences between different types of sources.

5 What is a Light Source?

In the last section, we assumed that a ray manifold could be partitioned into source and non-source rays, but we did not show how to do so. To appreciate the subtlety of this problem, consider an outdoor scene on an overcast day and let free space $\mathcal{F}$ be a finite volume within the scene. Most would agree that the overcast sky is a source, and so the set of source rays in $\mathcal{M}(\mathcal{F})$ would be the subset which, when extended backwards, would reach the sky.

A subtlety arises, however, when we consider whether a white piece of paper in this scene would also be 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 to be a source but a white piece of paper is not ?

One distinction between the sky and the piece of paper is that the rays from the sky transmit light into free space while the rays from the 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 this answer is unsatisfying. For example, a translucent object such as a 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 it transmits light into the room (a doorway leading outside would be a source, but a doorway leading to a closet would not be).

It seems that to partition a given ray manifold into source and non-source rays it is not enough to know the emitted, transmitted, and reflected components of each ray. It seems we must also know the physics of what is beyond the boundary of free space, *i.e.* the sun behind the cloud, and the cloud beyond the doorway.

But again this is unsatisfying. First, since the radiance in a given free space is entirely determined by the physics at the boundary, it seems that there should exist a physics–based definition that doesn't rely on the physics beyond the boundary. Second, by relying on the physics beyond the boundary, we do not gain insight regarding vision, since what is happening beyond the boundary is often not visible from points in free space. We want eventually to understand how lighting

Table 1: The sixteen corners of the 4-D light source hypercube. These corners are defined by taking the parameters, h_x, h_y, h_p, h_q, to their limits, *i.e.* 0 or ∞.

Non-ideal example	*Ideal model*	h_x	h_y	h_p	h_q	dimension
overcast sky	uniform source	∞	∞	∞	∞	4
Cyberware^{TM} scanner		∞	∞	∞	0	3
		∞	∞	0	∞	
fluorescent tube	linear source	∞	0	∞	∞	3
		0	∞	∞	∞	
sunlight	point source at infinity	∞	∞	0	0	2
	uniform distribution of rays in a plane	∞	0	∞	0	2
		0	∞	0	∞	
louvered linear source (see text)	fan of rays perpendicular to a linear source	∞	0	0	∞	2
		0	∞	∞	0	
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	
rotating spotlight	fan of rays	0	0	0	∞	1
		0	0	∞	0	
spotlight or laser	single ray	0	0	0	0	0

conditions can be inferred by a vision system that is within a scene, not outside of it.

With this 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 this, we introduce a "thought experiment." Consider a demon that absorbs completely the light arriving at a given region of $\partial\mathcal{F}$ from outside of $\mathcal{F}$. Specifically, a demon absorbs the components of radiance that are emitted and transmitted into $\mathcal{F}$ along a single ray of $\mathcal{M}(\mathcal{F})$. 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, the set of source rays in $\mathcal{M}(\mathcal{F})$ is 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 "minimal" relation is defined by set inclusion on $\mathcal{M}(\mathcal{F})$.)

Definition 3 (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 the above definition of a set of source rays captures our intuition of why, in this example, the closet and vase would not be considered sources.

One final point. 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 (see [6, 7]). 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. Since these two sets of rays are topologically disconnected, it makes intuitive sense to think of them as distinct subsources. A more subtle example is a single window transmitting moonlight and also light from a street lamp. Although all source rays originate from the window, the (p, q) directions of the rays form two mutually–disjoint sets, and for this reason two subsources are defined. This example emphasizes why all four dimensions of the ray manifold are needed to articulate source composition.

6 An Example

Figure 2 shows a room illuminated by six different types of light sources: a 0-D a 1-D, two 2-Ds, a 3-D, and a 4-D source. The images were rendered using the RADIANCE lighting system [12]. (All surfaces in the scene are Lambertian, and interreflections are computed up to four bounces. Image intensities are thresholded to a maximum grey level value of 255.)

In each image, the source rays originate at the far left wall of the room. The radiance of each source is normalized such that each source has the same radiant flux (W), *i.e.* each source provides the same amount of light to the room (recall Sec. 4.1). Let us discuss each of the six sources shown.

The 0-D source is a narrow beam of rays, for example, sunlight passing through a small square hole in the left wall. (Note that the hole itself appears dark in the image since the sun is not visible through it.) The source acts as a spotlight, illuminating a rectangular region on the right wall. This rectangular region in turn illuminates the room via interreflections.

In rendering the scene, we found we could reduce artifacts by simulating the 0-D source with a 2-D source. To do this, we placed a rectangular 2-D light source (or panel light) of the appropriate shape and radiance directly on the right wall. The simulation is valid as long as the right wall has Lambertian reflectance, *i.e.* all rays from a given point on the wall have the same radiance.

Two observations follow. First, sources of different dimension can produce equivalent illumination. When a 0-D source illuminates a Lambertian surface, there are two consistent interpretations: a spotlight, and a 2-D point source (the spot). Note that this two-fold ambiguity occurs only when the scene is static. In a dynamic scene, a moving object passing through the 0-D source beam would become illuminated, and the bright spot on the right wall would disappear.

The second observation is that the dimensionality of the source determines the range of radiances in a scene. This is significant because an excessive range can lead to *glare*. (Glare is not apparent in Figure 2 since the grey levels intensities have been thresholded to a maximum value.) Real vision systems need to cope with glare frequently. Moreover, the way that a vision system can avoid glare depends on the dimensionality of the source. For example, one could avoid the glare from a 0-D or 1-D source by moving perpendicularly to it, whereas a different strategy would be required for a 2-D source. Clearly the problem of how to avoid glare is important, and the dimensional analysis of this paper provides constraint for addressing this problem.

We next consider the 1-D source which is defined by extending vertically the hole in the left wall. An example is sunlight passing through a crack in a doorway. As in the 0-D case, note that the brightest surface in the room is not the source, but rather the surface directly illuminated by the source, in this case the streak on the floor. (We have again simulated the source using a rectangular diffuse source. In this case, we simulate a 1-D source using a 3-D source.)

Two 2-D sources are shown in the middle row of Figure 2: a large doorway transmitting parallel rays of sunlight, and a small panel light emitting rays in all directions. Observe the shadows under the table. In each image, the table legs and the horizontal beams both produce sharp shadows.

A 3-D source is shown in the bottom row. An example is a crack in a doorway on a cloudy day. Note that the table legs cast sharp shadows but that the horizontal beams do not. The reason is that each table leg is roughly co-planar to the source and hence either occludes the source completely or not at all. The horizontal supports are perpendicular to the source and hence can occlude at most a small part of the source.

Finally, consider a 4-D source such as a large rectangular panel light or a large doorway transmitting the diffuse light from an overcast sky. Note that the contrast of the image is relatively small. This corresponds to the intuitive idea that cloudy days are "grey days." In particular, 4-D sources cannot produce sharp shadows.

7 Conclusion

The appearance of a scene depends crucially on how the scene is illuminated. In particular, to correctly interpret shading and shadowing cues, a vision system must successfully account for the lighting. Shadows can be smooth or sharp, and this smoothness or sharpness in general depends on position in the scene, the orientation of the object boundary casting the shadow, and especially on the dimension of the source.

Unfortunately, there presently exist no methods in computer vision for inferring the type of the light

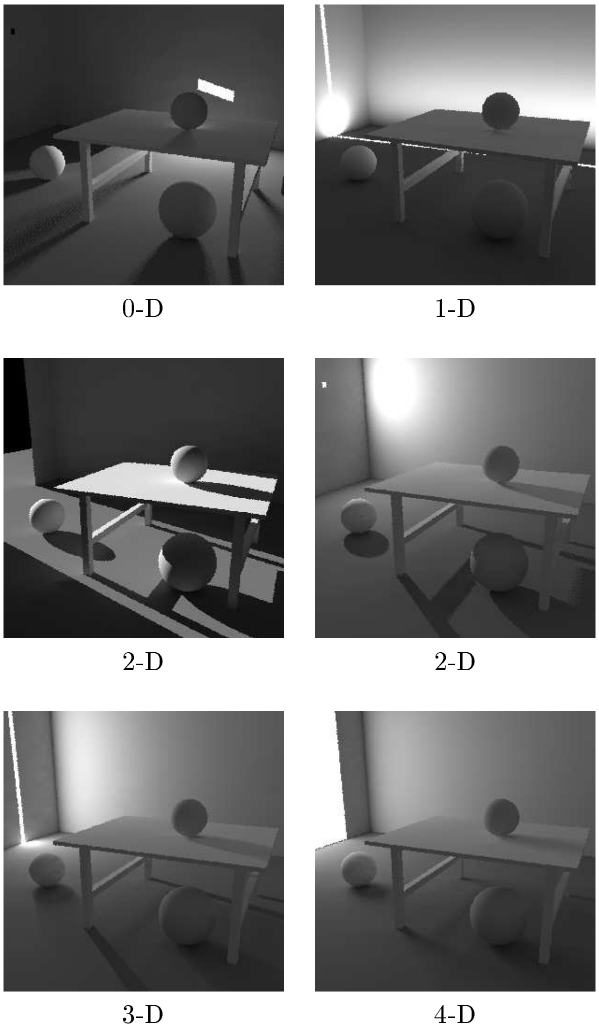

Figure 2: A scene illuminated by six different types of light sources.

source from an image. We regard this as a key limitation especially since one of the goals of computer vision is to build systems that can operate in non–contrived environments. Specifically, the problems of how a vision system can deal with glare, and how glare and image contrast can be used to infer the type of the source(s) have simply not been addressed.

We regard our contribution in this paper to be a basic one. We have developed a framework for defining and relating different types of light sources based on a dimensional analysis of the ray manifold. The framework specifies what is to be computed in solving the general light source estimation problem, namely, what are the dimensions of the source? and how is light distributed across these dimensions? Although initially our analysis may seem hopelessly abstract, we argue that it opens the door to constraints involving contrast, shadow sharpness, and the basic definition of a light source. Such constraints provide a bridge linking the problem of light source estimation to those of current mainstream vision research.

Acknowledgments

This research was supported in part by grants from NSERC and AFOSR, while the authors were at the Center for Intelligent Machines, McGill University. The authors would like to thank James Elder, David Jacobs, Allan Jepson, David Kriegman, John Oliensis, Kaleem Siddiqui, and Lance Williams for their comments.

References

[1] P. Belhumeur and D. Kriegman. What is the set of images of an object under all possible lighting conditions ? In *CVPR*, San Francisco, 1996.

[2] B. K. P. Horn. Understanding image intensities. *Artificial Intelligence*, 8:201–231, 1977.

[3] B. K. P. Horn and M. J. Brooks, editors. *Shape from Shading*. MIT Press, Cambridge, Mass, 1989.

[4] K. Ikeuchi. Determining surface orientations of specular surfaces by using the photometric stereo method. *PAMI*, 3(6):661–669, 1981.

[5] M. S. Langer and S. W. Zucker. Shape-from-shading on a cloudy day. *JOSA*, 11(2):467–478, 1994.

[6] M. S. Langer and S. W. Zucker. A ray-based computational model of light sources and illumination. In *IEEE Workshop on Physics-based Modelling in Computer Vision*, Cambridge, Ma., June 1995.

[7] M. S. Langer and S. W. Zucker. Casting light on illumination: a computational model and dimension analysis of sources. *Computer Vision and Image Understanding*, 1997. special issue on "Physics-based modelling and reasoning in computer vision" (in press).

[8] P. H. Moon. *The Scientific Basis of Illuminating Engineering*. Dover, New York, 1961.

[9] A. Pentland. Finding the illuminant direction. *JOSA*, 72:448–455, 1982.

[10] M. Planck. *The theory of heat radiation*. Dover, 1959.

[11] H. Stark and F. B. Tuteur. *Modern Electrical Communications*. Prentice Hall, Englewood Cliffs, NJ, 1979.

[12] G. Ward. The radiance lighting simulation and rendering system. *Computer Graphics*, pages 459–471, July 1994.