Illumination

Spatially Varying Illumination : A Computational Model of Converging and Diverging Sources

M. S. Langer and S.W. Zucker

Research Center for Intelligent Machines
McGill University
3480 University St. Montreal, H3A2A7, Canada
email: langer@cim.mcgill.ca, zucker@cim.mcgill.ca

Abstract. There are three reasons for illumination to vary within a scene. First, a light source may be visible from some surfaces but not from others. Second, because of linear perspective, the shape and size of a finite source may be different when viewed from different points in a scene. Third, the brightness of a source may be non-uniform. These variations are captured by a new computational model of spatially varying illumination. Two types of source are described: a distant hemispheric source such as the sky in which light converges onto a scene, and a proximal source such as a lamp in which light diverges into a scene. Either type of source may have a non-uniform brightness function. We show how to render surfaces using this model, and how to compute shape from shading under it.

1 Introduction

There are three reasons why illumination may vary within a scene. The first is that the source may be visible from some surfaces but not from others. The second is that, because of linear perspective, the shape and size of the source may be different when viewed from different points in the scene. The third is that the brightness of the source may be non-uniform.

In order to draw inferences about a scene from such illumination variations, we develop a model of how light flows through a scene. This model is computational in that it specifies data structures for representing general types of illumination variation, as well as algorithms for manipulating these data structures. Specifically, we model two general illumination scenarios: an outdoor scene illuminated by the sky, and a scene illuminated by a proximal diverging source such as a lamp or window. The model generalizes our earlier papers [1, 2] in which we assumed that the light source was a uniformly bright sky.

2 Visibility Fields

We begin with an example that illustrates the fundamental issues. Consider an empty room with dark colored walls (allowing us to ignore surface interreflections). The room is illuminated by light from the sky which passes through a

* This research was supported by grants from NSERC and AFOSR.

Lecture Notes in Computer Science, Vol. 801
Jan-Olof Eklundh (Ed.)
Computer Vision – ECCV '94
© Springer-Verlag Berlin Heidelberg 1994

window. Observe that *the fundamental cause of the illumination variation is that the set of directions in which the light source is visible varies across free space.* This geometric variation is explicitly shown in Figure 1. Each disc in this figure represents the set of light rays passing through a point in free space within the room. The white sectors of each disc represent the rays which come directly from the window (and hence from the source), and the grey sectors represent the rays which come directly from the walls.

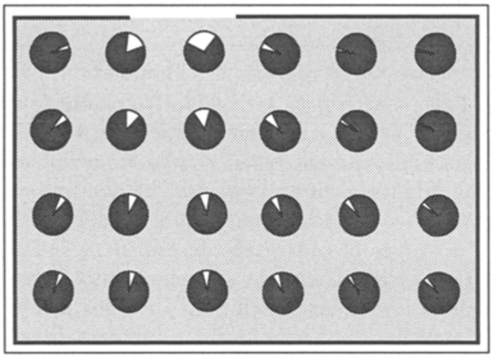

Fig. 1. An empty room illuminated by the light passing through a window.

This partition of the rays is the key tool for modelling illumination variation within a scene, and is formally defined as follows [1]. The VISIBILITY FIELD at a point $\mathbf{x}$ in free space is a set of unit length vectors, $\mathcal{V}(\mathbf{x})$, that represent the directions in which the light source is visible from $\mathbf{x}$. Observe that, as one moves away from the window in a particular direction, the solid angle of rays in which the window is visible decreases.

3 Brightness of a Light Ray

Visibility and aperture are geometric properties of free space. As such, they cannot account for photometric properties of light, in particular, that not all light rays carry the same energy density. For example, on a sunny day, certain of the light rays which pass through a window come from the sun, while others come from the blue sky.

Energy density of light is specified in terms of the brightness of a light ray, which is defined as follows [3, 5]. Consider viewing a scene through a narrow straight tube. Suppose that the near end of the tube is positioned at a point $\mathbf{x}$, and that the tube is pointing in direction $\mathbf{L}$. Let the light energy passing through the tube be d^2E. Let the solid angle subtended by the far end of the tube (when viewed from $\mathbf{x}$) be $d\Omega$. Let the cross sectional area of the tube be $d\mathbf{a}$. Then, the

BRIGHTNESS OF A LIGHT RAY passing through $\mathbf{x}$ from direction $\mathbf{L}$ is

$$B(\mathbf{x}, \mathbf{L}) \equiv \frac{d^2 E}{da\, d\Omega} .$$

Brightness has units lumens per square metre per steradian. The key property of the brightness is that, in the absence of scattering (eg. by fog), brightness is constant along a ray [3]. Considering each ray as a single geometric object, we may thus assign to this object a single brightness value.

For a light ray that originates from a reflecting surface rather than from a source, brightness may be associated with a surface point $\mathbf{x}$ and with the direction $\mathbf{L}$. This brightness value is referred to as the LUMINANCE of $\mathbf{x}$ in direction $\mathbf{L}$. In the case of a Lambertian surface, luminance is independent of direction, and may be modelled as

$$B_{out}(\mathbf{x}) = \frac{\rho}{\pi} \int_{\mathcal{V}(\mathbf{x})} B(\mathbf{x}, \mathbf{L})\, \mathbf{N}(\mathbf{x}) \cdot \mathbf{L}\, d\Omega . \tag{1}$$

This model is accurate provided that the albedo is low, that is, provided that surface interreflections may be ignored.

4 Converging and Diverging Sources

Two extreme scenarios are of special interest. The first is a light source that is much larger than a scene, so that rays from the source *converge* to the scene. The canonical example is the sky. The second is a light source that is much smaller than a scene so that rays from the source *diverge* to the scene. The canonical example is a lamp or candle.

The brightness of rays coming from a converging or a diverging source depend only on direction. This is obvious for a converging source, since the brightness of the sky does not vary as one moves within a scene. The case of the diverging source is less obvious, but may be understood in terms of an example of a room illuminated by sky light that passes through a clean window. Light rays that converge on the window diverge into the room, so that the rays entering the room inherit the brightness of rays from the sky, which depend on direction only.

Let $B_{src}(\mathbf{L})$ denote the brightness function of a converging or diverging source, and write (1) as

$$B_{out}(\mathbf{x}) = \frac{\rho}{\pi} \int_{\mathcal{V}(\mathbf{x})} B_{src}(\mathbf{L})\, \mathbf{N}(\mathbf{x}) \cdot \mathbf{L}\, d\Omega . \tag{2}$$

It is important to note that in classical photometry, a diverging source is usually approximated as a point[5] whose illumination is specified by *luminous intensity* which has units lumens per steradian. This definition, however, does not allow certain causes of illumination variation to be distinguished. For example, a source may have non-spherical shape or it may have non-uniform brightness, or both. Moreover, luminous intensity does not account for the penumbra (smooth cast shadow boundaries) produced by the finite size of the source. The advantage of the above model (2) is that it explicitly models the actual causes of illumination variation in a scene.

5 Discretization of Model

In order to perform computations, the variables of our models must be discretized. Space may be represented by an $N \times N \times N$ cubic lattice. A node in this lattice is $\mathbf{x} = (x, y, z)$. Assume that a surface seen in an image may be represented by a continuous depth map $\hat{z}(x, y)$ defined on a unit square. Given such a surface, a set of free space nodes $\mathcal{F}$ is the set of nodes lying above the surface.

Light travels from one free space node to another. Light is restricted to travel in a small number of directions, which are defined by an $M \times M$ cube, where $M \ll N$. Each node on the surface of this cube defines a direction, namely, the direction of a light ray passing through that node and through the center of the cube. In particular, we restrict our discussion to sources that are on the same side of the plane $z = 0$ as the viewer, so that the directions of light rays coming from the source may be represented by a hemicube [4], denoted $\mathcal{H}^*$.

A brightness function, $B^*_{src}(\mathbf{L})$, is defined on $\mathcal{H}^*$. Because the directions $\mathcal{H}^*$ define a non-uniform spacing of the unit sphere, to discretize the integral it is necessary to weight each of directions by a solid angle $\Delta\mathbf{L}$ that depends on $\mathbf{L} \in \mathcal{H}^*$. Finally, for each node $\mathbf{x}$, let the discrete set of directions in which the light source is visible from $\mathbf{x}$ be denoted $\mathcal{V}^*(\mathbf{x})$, so that $\mathcal{V}^*(\mathbf{x}) \subseteq \mathcal{H}^*$ Then, for either a converging or diverging source, we have the following model of surface luminance,

$$B_{out}(\mathbf{x}) = \frac{\rho}{\pi} \sum_{\mathbf{L} \in \mathcal{V}^*(\mathbf{x})} B^*_{src}(\mathbf{L})\, \mathbf{N}(\mathbf{x}) \cdot \mathbf{L}\, \Delta\mathbf{L}\,. \qquad (3)$$

In the next section, we present computational algorithms which are based on this model.

6 Forward and Inverse Algorithm

We present a generalization of the forward and inverse algorithms introduced in [1], where a uniform converging source was assumed. We now consider both converging and diverging sources, each having arbitrary brightness functions.

Light rays enter a scene through the boundary of free space $\mathcal{F}$. For a given scene, the set of rays coming from a source may be specified by the values of the visibility field on the boundary of free space. Both the forward and inverse algorithms depend on the computation of the visibility field over free space, $\mathcal{F}$. This computation is performed by propagating the visibility field away from the boundary of $\mathcal{F}$. The computation is by induction. The boundary condition on the visibility field is given for depth $n = 0$. Then, assuming the visibility field has been computed up to depth n, it is computed at depth $n + 1$.

FORWARD ALGORITHM: Given $z(x, y)$, compute $B^*_{out}(x, y)$.

```
n := 0 ;
repeat
   for all (x, y),
      x := (x, y, n);
      if n ≤ z(x, y)
         then for all L ∈ H*
              if L ∈ V*(x + L) and x + L ∈ F
              then L ∈ V*(x)
              else L ∉ V*(x);
      if n = z(x, y), then compute B*_out(x, y) using Eq. (3);
   n := n + 1;
until for all (x, y),  z(x, y) < n .
```

The inverse problem is more challenging. As in [1, 2], we ignore the shading effects of the surface normal. This is done by replacing the factor $\mathbf{N}(\mathbf{x}) \cdot \mathbf{L}$ by its average value, 0.5, on the unit hemisphere $\mathcal{H}(\mathbf{x})$, yielding the model

$$B^*_{out}(\mathbf{x}) \equiv \frac{\rho}{2\pi} \sum_{\mathbf{L} \in \mathcal{V}^*(\mathbf{x})} B^*_{src}(\mathbf{L}) \ \Delta\mathbf{L} \ . \tag{4}$$

The present model is more general since it allows for directional variation in the brightness of the source.

INVERSE ALGORITHM: Given an image $I(x, y)$, compute $z(x, y)$.

```
for all (x, y),  z(x, y) := 0 ;
n := 0 ;
repeat
   for all (x, y) ∈ P*,
      x := (x, y, n);
      if z*(x, y) = n
         for all L ∈ H*
            if L ∈ V*(x + L) and x + L ∈ F*
            then L ∈ V*(x)
            else L ∉ V*(x);
         Compute B*_out(x) using Eq. (4);
      if | B*_out(x) − I(x, y) | > ε then z(x, y) := n + 1;
   n := n + 1;
until for all (x, y),  z(x, y) < n .
```

7 Results

A slanting plane is rendered using a diverging spherical source. Three brightness functions are used: isotropic (left), weakly directed (middle), and strongly directed (right). As the brightness becomes more directed along the optical axis, the maximum of the image intensity shifts toward the center of the image.

A depth map is computed from each of the three images. The corners of the example on the right illustrate an important ambiguity in the inverse algorithm. Since a spotlight gives off a cone of light, the colums 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.

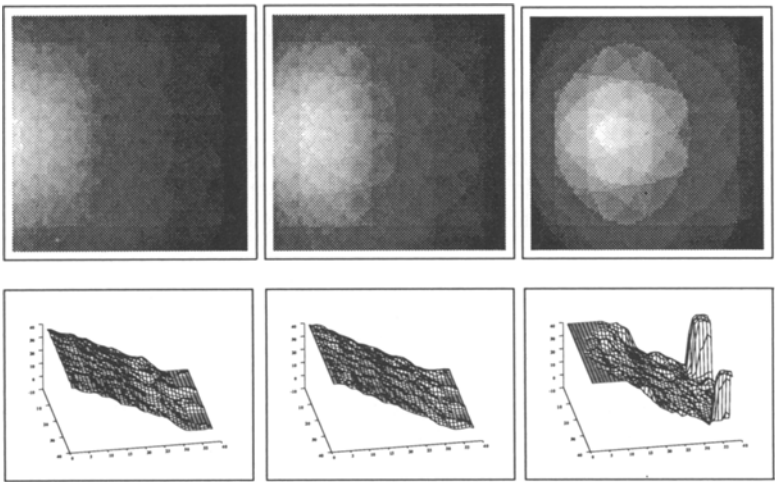

Fig. 2. The diverging source is a sphere centered at the viewer.

References

1. M. S. Langer, S.W. Zucker, "Diffuse Shading, Visibility Fields, and the Geometry of Ambient Light", *Proc. Fourth ICCV*, Berlin, Germany. May 1993.
2. M. S. Langer, S.W. Zucker, "Shape from Shading on a Cloudy Day". *J. Opt. Soc. Am.* (in press).
3. A. Gershun, "The Light Field", J.Math.Phys. 18,51-151 (1939).
4. Cohen, M.F., Greenberg, D.P. "The hemicube: A radiosity approach for complex enviroments." Computer Graphics 22(4) 155-164 (1985).
5. P. Moon, The Scientific Basis of Illuminating Engineering. (Dover Publications, 1961).