Mutual illumination

David Forsyth and Andrew Zisserman *
Robotics Research Group
Department of Engineering Science
Oxford University
England

Abstract

Although mutual illumination has been largely ignored by the vision community to date, its effects are easily observed in images. Mutual illumination effects can introduce serious errors into any shape from shading scheme based on the image irradiance equation, because in the presence of mutual illumination, radiance is no longer solely dependent on surface orientation. However, as the magnitude of these effects is strongly related to the albedo of surfaces and to surface geometry, they can provide useful cues to absolute surface lightness and to the arrangement of surfaces.

We report here some theoretical and experimental results which underline the importance of mutual illumination to visual modules dealing with shape and with surface lightness. Our experiments are in good agreement with results obtained with a simple theoretical model. These results show the effects of mutual illumination in pictures of simple objects, and indicate that these effects must be accounted for in modelling image intensities.

1 Introduction

When one illuminates a world containing any more complex arrangement of objects than a single convex surface, one discovers that light reaches a surface facet from many sources. Objects reflect light not just to sensors, but also onto other objects. This leads to a pattern of measured radiances that can differ significantly from those predicted by models which assume that radiance is purely a property of the orientation of a surface facet relative to the light source.

A large number of techniques have been proposed for recovering surface shape from shading patterns (see Horn's book [6] for a clear discussion of this field). Such techniques are forced to ignore the effects of mutual illumination, because the global nature of these effects makes it difficult to formulate recovering local shape from radiance information in their presence. Accurately simulating mutual illumination is, however, recognised to be important in obtaining realistic computer graphics [3], and the ability to predict mutual illumination patterns is valuable to photometrists [8].

The basic equation governing mutual illumination in a scene is a Fredholm equation of the second kind. It is often difficult to deal with such equations intuitively, and so we supply here a description of the meaning and display the form of solutions of this equation for some simple geometries.

Our experimental work indicates that one dimensional approximations can be applied successfully in situations which have a local translational symmetry. This work also shows that mutual illumination effects can not only offer valuable cues to surface lightness perception, but also confound shape recovery.

In the body of this paper, we discuss theoretical and numerical approaches to predicting the effects of mutual illumination. We show results of these approaches for a number of simple geometries, and show experimental results obtained from images of real objects. These results allow us to speculate on the possible value of mutual illumination effects to visual processes that recover surface shape and surface lightness.

*DAF acknowledges the support of the Rhodes Trust. AZ acknowledges the support of the Science and Engineering Research Council.

2 Theory

2.1 The mutual illumination equation

We use the photometric terms **radiance**, **irradiance** and **albedo** here, and explain these terms below:

Irradiance

Irradiance refers to the amount of light falling on a surface, and is conventionally measured in Watts per square metre.

Radiance

Radiance refers to the amount of light reflected from a surface. This is in most cases equivalent to the "brightness" of the surface, and is measured in Watts per square metre per unit solid angle.

Albedo

Albedo refers to the fraction of the light incident on a surface that is in fact reflected.

Consider a scene consisting of surfaces parametrised in some way that allows us to write the surfaces as $\mathbf{r}(\mathbf{u})$, where the bold font denotes a vector quantity. At a point $\mathbf{u}$, we denote the surface normal by $\mathbf{n}(\mathbf{u})$. Introduce the term $\mathbf{d_{uv}}$ for $\mathbf{r}(\mathbf{u}) - \mathbf{r}(\mathbf{v})$, and assume that all surfaces are Lambertian. We write the radiance at the point parametrised by $\mathbf{v}$ as $N(\mathbf{v})$. The radiance ("brightness") $N(\mathbf{v})$ arising from $N(\mathbf{u})$ is:

$$\frac{\rho(\mathbf{v})}{\pi} \frac{< \mathbf{n}(\mathbf{v}), \mathbf{d_{uv}} >< \mathbf{n}(\mathbf{u}), \mathbf{d_{vu}} >}{< \mathbf{d_{uv}}, \mathbf{d_{uv}} >^2} View(\mathbf{u}, \mathbf{v})\ N(\mathbf{u})$$

Where $View(\mathbf{u}, \mathbf{v})$ is 1 if these points have a line of sight from one to the other, and 0 otherwise, $N(\mathbf{u})$ is the radiance at $\mathbf{u}$, $<,>$ denotes the dot product operation on two vectors, and $\rho(\mathbf{v})$ denotes the albedo of the surface patch parametrised by $\mathbf{v}$. The terms of the form $< \mathbf{n}(\mathbf{u}), \mathbf{d_{uv}} > / < \mathbf{d_{uv}}, \mathbf{d_{uv}} >^{\frac{1}{2}}$ express cosines for Lambertian reflection, and the term $1/ < \mathbf{d_{uv}}, \mathbf{d_{uv}} >$ expresses inverse square law effects. Informally, the term $\rho(\mathbf{v})/\pi$ serves to turn incident light into radiance.

Since we know that the radiance at a point is the sum of that due to the radiance at all other points and that due to the light source, we obtain the following equation for the radiance at $\mathbf{u}$:

$$N(\mathbf{u}) = N_0(\mathbf{u}) + \rho(\mathbf{v}) \int_D K(\mathbf{u}, \mathbf{v}) N(\mathbf{v}) d\mathbf{v} \qquad (1)$$

where

$$K(\mathbf{u}, \mathbf{v}) = \frac{1}{\pi} \frac{< \mathbf{n}(\mathbf{v}), \mathbf{d_{uv}} >< \mathbf{n}(\mathbf{u}), \mathbf{d_{vu}} >}{< \mathbf{d_{uv}} >^2} View(\mathbf{u}, \mathbf{v})$$

Where D refers to the whole domain of the parametrisation, and $N_0(\mathbf{v})$ is the component of radiance at $\mathbf{v}$ due to the effects of the source alone. For example, for the simple case of an effective point light source at infinity, with direction $\mathbf{s}$ and brightness B, we have $N_0(\mathbf{v}) = \rho(\mathbf{v})\ B < \mathbf{s}, \mathbf{n}(\mathbf{v}) >$. We call equation 1, which expresses energy balance, the **mutual illumination equation**. Solving this equation analytically is simplified by assuming that $\rho(\mathbf{v})$ is constant. Equation 1 then becomes a Fredholm equation of the second kind, and K is referred to as the **kernel** of this equation. It is easy to see that

for physical reasons,

$$\int_D \int_D K^2(\mathbf{u},\mathbf{v}) d\mathbf{u} d\mathbf{v} \leq 1$$

for if this were not true, some surface would be contributing extra energy.

2.2 Analytic solutions to the mutual illumination equation

There are several approaches for solving equations of this type (see, for example, [9]). One well-known solution is given by a Neumann series:

$$N(\mathbf{u}) = N_0(\mathbf{u}) + \sum_{n=1}^{n=\infty} \rho^n \int_D K_n(\mathbf{u},\mathbf{v}) N_0(\mathbf{v}) d\mathbf{v}$$

Where

$$K_n(\mathbf{u},\mathbf{v}) = \int_D K_{n-h}(\mathbf{u},\mathbf{w}) K_h(\mathbf{w},\mathbf{v}) d\mathbf{w}$$

and

$$K_1 = K,\ n \geq 2,\ 1 \leq h \leq n-1\ )$$

Ample physical evidence exists that this series converges (in fact, to prove this for $\rho < 1$, notice this series converges for $|\rho| < \| K \|^{-1}$ ([9], p.50 *et. seq.*), and that $\| K \|^{-1} \geq 1 > \rho$).

When ρ is small, the solution becomes a solution of the image irradiance equation as the effects of mutual illumination will be dominated by source effects. Thus, to observe experimentally the effects of mutual illumination, we need two set of objects: one white, and the other black. Qualitative differences in radiance distributions for images taken of similar arrangements of these objects can then be ascribed to the effects of mutual illumination.

Although we know the Neumann series will converge, it is not always possible to say very much about the **rate** of convergence of the series. The series has the following physical meaning:

> To form the radiance at any point, take the radiance at that point due to the effects of the source. Now compute the effects due to rays that are reflected once to arrive at that point, and add them to this radiance. Now compute the effects of rays that are reflected twice, and add them. Continue computing the effects of rays that are reflected n times untill n is very large. This procedure is often referred to as "ray tracing".

The difficulty with this approach is that one may think up objects sufficiently contorted in shape that rays reflected n times have no significant effect, but those reflected $n+1$ times do. We show a simple example in figure 1 (Photometrists call this effect vignetting). This point is important only for objects with high albedoes, for the effects of these rays will otherwise peter out rather quickly.

Equation 1 has another useful property, best seen by inspection of the Neumann series solution. If we consider the Neumann series as an operator on (integrable!) initial radiance functions (written above as $N_0(\mathbf{u})$), it is **linear**. Loosely, if you want to know the radiance you would get with two light sources, you may measure radiances under both separately, and add the results. This is tremendously valuable from an applications point of view. If we believe that there is a limited set of types of illumination, for example, that all scenes are illuminated by a weighted sum of an isotropic ambient illumination and an effective point source at infinity, we know that this property guarantees that we need only work out the mutual illumination effects for the two sources separately, to know the form of the effect for the sum of the sources.

The operator that takes the initial radiance distribution to the radiance distribution observed is, however, a complicated function of global scene geometry. We can nonetheless claim that this linearity property means that if we keep on hand a dictionary of the effects of a small number of different lighting distributions on a number of different geometries we are in a strong position for analysing the effects of mutual illumination in images, because we may compute the effects of any weighted sum of these lights on these geometries.

Another standard method for solving Fredholm equations of the second kind, involves expanding $N_0(\mathbf{v})$ on a basis consisting of the eigenfunctions of the kernel. This leads to an expansion of $N(\mathbf{v})$ on the same basis. This approach is particularly attractive when the kernel is self-adjoint (or, for real kernels, symmetric), as the eigenfunctions are guaranteed to be orthogonal. This approach is used by Koenderink and Van Doorn [8], who call the eigenfunctions "pseudo facets", referring to the fact that by employing this expansion one may rewrite the mutual illumination equation as a form of the image irradiance equation for a surface with a different geometry. It is a particularly valuable approach when the eigenvalues diminish quickly, as in this case only the terms corresponding to a small number of the largest eigenvalues are necessary to obtain a good approximation to the solution.

2.3 Interpreting mutual illumination information

The kernel of equation 1 contains a tremendous amount of global geometric information. Our experimental data suggests that it may be possible to model the contribution of "far-away" surfaces as isotropic ambient illumination, but that the effects of nearby surfaces are important and are anisotropic. Part of our ongoing research involves extracting geometric information from the kernel of this equation, employing one's knowledge of its solution, that is, the observed image intensities. However, it is easy to see that the effects of mutual illumination will be most prominent in regions where surfaces face one another, and where they are close, for example, for concave surface patches, or in those places where surfaces meet to form a concave region.

The magnitude of local effects taken with the magnitude of the ambient illumination field, provide strong cues for absolute surface lightness information. At least qualitative absolute lightness information may be extracted, given that we have a reliable source of shape information (for example, stereopsis) by comparing the radiance distribution resulting from the known shape with that observed. An implementation of this approach would be a major step forward for those studying surface lightness perception.

We present an analytic solution to the mutual illumination equation for a simple geometry. This analytic solution indicates the type of geometric constraints available when only low order terms of the Neumann series contribute to the solution. For more complex geometries, we use numerical solutions obtained using a finite element method. These solutions have been constructed for geometries with a translational symmetry. Infinite objects are not easy to simulate experimentally, so that translational symmetries were purely local in the experiments. However, comparison with experimental results suggests that using an assumption of translational symmetry does not prevent a good approximation to observed data. This suggests that in many cases, a one parameter model is appropriate.

2.4 Solutions for a single translational symmetry

2.4.1 The form of the kernel for a single translational symmetry

We consider a coordinate system with translational symmetry along the y axis, depicted in figure 2. Thus, all relevant shape information is contained in any plane $y = constant$. Consider the plane $y = 0$. Now for an infinitesimal segment on the boundary of the shape in this plane, it is clearly the case that we may construct a 2 dimensional kernel by integrating along the y axis the contribution of the surface facets that project to this segment

Then consider $\mathbf{r}(\mathbf{u})$, $\mathbf{r}(\mathbf{v})$, which are three-vectors lying on the surface. In particular, we have the unique decomposition:

$$\mathbf{r}(\mathbf{u}) = \mathbf{x}(s) + \mathbf{j}.y$$

For some appropriate reparametrisation, with $\mathbf{x}$ restricted to the plane $y = 0$. We now have, for K^* the 2-dimensional version of the kernel:

$$K^*(s,s^*) = \int_{-\infty}^{\infty} K dy$$

We can parametrise the 2D kernel in terms of 2 arc length parameters s and s^*, because all shape information is encapsulated in a set of plane curves. In particular the normals to surface facets must be orthogonal to the direction of translational symmetry. Thus we obtain:

$$K^*(s,s^*) = \frac{1}{\pi} < \mathbf{n}(s), \mathbf{x}(s^*) - \mathbf{x}(s) > \ < \mathbf{n}(s^*), \mathbf{x}(s) - \mathbf{x}(s^*) > I(s,s^*)$$

where

$$I(s,s^*) = \int_{-\infty}^{\infty} \frac{dy}{(< \mathbf{x}(s^*) - \mathbf{x}(s), \mathbf{x}(s^*) - \mathbf{x}(s) > + y^2)^2}$$

Carrying out the integration gives us:

$$K^*(s,s^*) = \frac{1}{2} \frac{< \mathbf{n}(s), \mathbf{x}(s^*) - \mathbf{x}(s) > \ < \mathbf{n}(s^*), \mathbf{x}(s) - \mathbf{x}(s^*) >}{< \mathbf{x}(s^*) - \mathbf{x}(s), \mathbf{x}(s^*) - \mathbf{x}(s) >^{3/2}}$$

An important analytic solution to the 2 dimensional version mutual illumination equation is the solution in the case of two planar facets, which we solve below as an example.

2.4.2 Mutual illumination for two planar facets with a translational symmetry

Consider the geometry shown in figure 3. For this figure, we have:

$$K(s,s^*) = \frac{sin^2(\alpha)}{2} \frac{s.s^*}{(s^2 + 2s.s^* cos(\alpha) + s^{*2})^{\frac{3}{2}}}$$

Now, we take a Neumann series to order ρ^2, to find:

$$N(s) = N_0(a) + N_1(s) + N_2(s)$$

where

$$N_1(s) = \rho N_0(b) \frac{sin^2(\alpha)}{2} s \int_{L_{b0}}^{L_{b1}} \frac{s^*.ds^*}{(s^2 + 2s.s^* cos(\alpha) + s^{*2})^{\frac{3}{2}}}$$

$$N_2(s) = \rho^2 N_0(a) (\frac{sin^2(\alpha)}{2})^2 s \int_{L_{b0}}^{L_{b1}} \frac{s^{*2}}{(s^2 + 2s.s^* cos(\alpha) + s^{*2})^{\frac{3}{2}}} J(s^*) ds^*$$

where

$$J(s^*) = \int_{L_{a0}}^{L_{a1}} \frac{t^2}{(s^{*2} + 2s^*.t.cos(\alpha) + t^2)^{\frac{3}{2}}} dt$$

Furthermore, we can obtain the coefficient of the term of order ρ in the Neumann series in closed form as:

$$\int_{L_{b0}}^{L_{b1}} K(s,s^*) ds^* = \frac{1}{2} \left[\frac{s + s^* \cos(\alpha)}{(s^2 + s.s^*.\cos(\alpha) + s^{*2})^{\frac{1}{2}}} \right]_{L_{b0}}^{L_{b1}}$$

This important result finds further application in our finite element work.

In fact, as $L_a, L_b \rightarrow L \rightarrow \infty$, we can write the entire series for the image radiance in the closed form:

$$N(s) = \frac{N_0(a) + N_0(b).\rho.sin^2(\frac{\alpha}{2})}{(1 - \rho^2.sin^2(\frac{\alpha}{2}))^2}$$

Thus, each plane's measured brightness receives a constant boost. A similar result is obtained by Horn [5], despite the fact that he did not consider inverse square law effects.

2.4.3 Numerical solutions: the finite element method

We consider a curve Γ, representing a section of a geometry with translational symmetry. We can write equation 1 for this geometry as:

$$N(s) = N_0(s) + \rho \int_{\Gamma} K(s,s^*) N(s^*) ds^* \quad (2)$$

Now approximate Γ by a set of linear elements Γ_i, so that Γ is approximated by $\sum_i \Gamma_i$. Then equation 2 becomes:

$$N(s) = N_0(s) + \rho \sum_i \int_{\Gamma_i} K(s,s^*) N(s^*) ds^* \quad (3)$$

Assume $N(s)$ is constant in these elements. We consider $N(s_i)$ for a node at the centre of an element, a process known as **collocation**. Then the mutual illumination equation becomes:

$$N(s_i) = N_0(s_i) + \rho \sum_j N(s_j) \int_{\Gamma_j} K(s_i,s^*) ds^*$$

Call $\int_{\Gamma_j} K(s_i,s^*) ds^* = K_{ij}$. Since we know K_{ij} in closed form (see above: we used a linear approximation to Γ precisely because we have this result in closed form), we need not use numerical integration, and can generate the matrix with i,j'th component K_{ij}.

The advantages of a finite element method are considerable: we may in fact allow ρ to vary over the geometry, we may use any illumination, and we can quickly and easily model any geometry with a translational symmetry. It is also easy to model self occlusion effects.

To obtain a solution, we need to solve the matrix equation

$$(I - PK)N = N_0$$

For P a diagonal matrix expressing the spatial variation of ρ. We used Gaussian elimination to solve this equation for N.

3 Results

3.1 Preliminary information

All images were taken with a CCD camera, with its automatic gain control defeated. Objects were either of matte white paper, with no visible surface texture, or were painted with either matte white or matte black enamel spray paint. The ratio of reflectance for the white paper to paper painted black was at least 40:1, and the paper painted black looked deep black in bright light.

It is difficult in practice to provide a point source at infinity, without involving mutual illumination effects other than those one wishes to measure. We therefore used a diffuse light source and took control images at each stage to ensure that the illumination over the area viewed was near uniform. Unfortunately, this means that the assumption in the theoretical development that all surfaces, even those outside the picture, are evenly illuminated, does not accurately reflect the circumstances of these experiments. Nonetheless, our experimental results agree extremely well with those predicted by the theory.

The philosophy behind these experiments was to look for qualitative differences between pictures of black objects and pictures of white objects. The image intensity function in pictures of black objects will conform to the expectation that the only source of illumination is a single distant light, and the image intensity function in pictures of white objects will show as well the effects of mutual illumination. All graphs show image intensity plotted against distance along a section.

3.2 A corner

We imaged an intersection between two planar patches for different values of the corner angle. The intersection was viewed along the angle bisector. The illuminant direction for each figure is indicated in the figure captions. We exploit the translational symmetry in this configuration and show the images as sections across the grey level function, but stress that the original image was a proper picture. It can be seen quite clearly from figures 5-10 that mutual illumination causes a significant qualitative effect (similar to the well-known "roof edge") in these images. Clearly, shape estimation based on image irradiance alone will err grossly and qualitatively for these images unless the effects of mutual illumination are taken into account.

Horn [5] constructed a numerical solution to a similar equation for this geometry and obtained the typical roof-edge signature, and recognised this to be a cue for concave polyhedral edges. Brady [1] makes this point too, and demonstrates the output of a roof edge finder that found a concave intersection of planar patches in this way.

In fact, the typical signature is a roof edge, or, more accurately, a

pair of rather "peaky" reflexes only when the illuminant lies along the angle bisector. This is clearly shown in figures 9 and 10. When the illuminant lies off the angle bisector, these reflexes are superimposed on a step edge, as shown in figure 5 and figure 7. The size of these reflexes is a powerful cue for albedo. Figure 7 shows how these reflexes vary dramatically in size with albedo. This implies that they could be employed as a cue to absolute surface lightness and to the angle of the intersection, perhaps by matching the qualitative appearance of the reflex to a "dictionary" of known shapes.

Typically, human observers appear to see a light patch on one side of the intersection of two planes, and a dark patch on the other (look at the corner of a room: the effect is present in figure 4, but may be reduced by the reprographic process). These patches are relatively narrow in extent (of the order of small numbers of millimetres), and run parallel to the intersection. One cannot explain these patches as the result of the effects of mutual illumination alone, as it is clear from figure 7 and from figure 5 that mutual illumination causes **both** sides of the step edge to increase in brightness. Some form of contrast effect may, however, be responsible.

3.3 Concave cylinders on a black background

Figure 12 shows the radiance profile for a concave black cylinder on a black background. Figure 13 shows the radiance profile for a concave white cylinder on a white background. For this case, we need to consider both the radiance distribution due to illumination by a point source at infinity, and that due to isotropic ambient illumination[1]. Since the operator that takes illumination to observed radiance distribution is linear in illumination, we may simply consider the form of a weighted sum of these effects in evaluating our results.

Figure 15 shows the effects of mutual illumination on the radiance distribution for a concave cylinder, computed for a point source of illumination at infinity which points directly into the cylinder. In particular, notice that the radiance distribution flattens noticeably. Figure 16 shows the radiance distribution resulting when the only illuminant is isotropic ambient illumination, where the radiance of centre surface facets is lower than that of facets on the side of the cylinder. Both these effects may be qualitatively understood by considering the 1st term in the Neumann series. The surface facets near the patches marked A in figure 14 will "see" more of the other facets more efficiently than those near the patch marked B because of the cosine term in the kernel. In the case of isotropic ambient illumination, the initial radiance is the same for all facets, so that mutual illumination effects will make the central facets look darker than those at the edges, whereas in the case of a point source at infinity, the central facets remain brighter than the edge facets, but brighter to a lesser extent than a reflectance map model would predict. In this case, therefore, the first term of the Neumann series provides a good guide to the overall effects of mutual illumination.

Since the solution is linear in the light sources, we may claim that the distribution observed experimentally arises as a weighted sum of the effects, and it is easy to see from the figures that such a weighted sum gives qualitatively the same shape as the observed distribution,shown in figure 13. There remains a large qualitative difference between the radiance profiles observed for a white concave cylinder, and for a black one. Since this qualitative difference in the radiances is strongly coupled to albedo and to properties of the illumination, it seems possible to use the radiance either in concert with albedo information to correct shape estimates, or in concert with more accurate sources of shape information to recover albedo, using some form of "dictionary" structure. In fact, it can be shown that for a concave hemi-cylinder, mutual illumination effects are independent of the radius of the cylinder.

3.4 Blocks world

Recently, Gilchrist [4] proposed that mutual illumination effects help people achieve surface lightness judgements. This suggestion is based on an experiment where he prepared small rooms, one white, one black, and filled each with objects of the appropriate colour. In particular, no changes in the retinal image could be ascribed to changes in surface properties. These rooms were then shown under different lighting conditions to human subjects, who could make good reports of surface lightness, despite the apparent poverty of the stimulus. Gilchrist points out that there are enormous differences in the qualitative behaviour of sections of images of each room.

We have built two similar rooms and imaged them. Figure 17 shows a typical image of the black room, and figure 18 shows a typical image of the white room. Notice in particular that there are many reflexes in figure 18, whereas figure 17 shows no sign of the effects of mutual illumination at all.

Figure 20 shows a section of the image intensity for the white room. Compare this with figure 19 which shows a similar section for the black room. Notice how, although the average grey levels are roughly similar, figure 19 shows a very much greater dynamic range than does figure 20, an effect apparently due to the high level of fairly isotropic ambient illumination generated by the white surfaces in the image of the white room. Furthermore, figure 20 shows a number of sloping edges that cannot be predicted with the simplest irradiance model. However in figure 19, typically the grey levels change sharply, and are coupled to surface orientation in the obvious way.

3.5 Convex cylinders

In this set of experiments, we took a number of images of white convex half-cylinders on black and on white backgrounds. We show here the most interesting effects from this series. These we compare with results obtained by numerically solving the mutual illumination equation. The geometry and terminology that we use are illustrated in figure 21.

Figure 23 shows a section of the image intensity measured for a white cylinder on a black planar background, with the camera perpendicular to the background. The light is at approximately 45 degrees to the background. The image intensity function is what one would expect. Figure 24 shows what happens when we replace the black background with a white background. Notice in particular the loss of dynamic range, and the large reflex where the cylinder meets the plane.

Figure 24 shows another interesting effect that one obtains with this geometry. Consider the section of the cylinder that is self-shadowed; all the surface facets in this section see the entire background plane, whereas the facets in the cylinder's cast shadow see only the shadowed section of the cylinder. As a result, we obtain a small but significant reflex in the self-shadow area. This reflex is marked in figure 24. This reflex raises a number of interesting possibilities. Consider, for the moment, this geometry illuminated only by a point source at infinity. In fact, as the self shadow region is not illuminated, any radiance here is due to mutual illumination, and as a result there are only odd order terms in the Neumann series. Under these circumstances and for a sufficiently large background plane, this reflex is an effect of the first order term in the Neumann series. In this case, the size of the background plane means that, intuitively, it radiates a lot of light onto the self shadow region. Third order effects must come from this light, reflected into the cast shadow region, and re-reflected, and are insignificant. As a result the effect of the background plane dominates all other effects. This reflex must occur in this geometry in all cases where there is a cast shadow which is sufficiently smaller than the background.

It is therefore possible to write down a simple equation for the reflex in terms of the extent of the plane, the lightness of the surfaces, and the orientation of surface facets as a function of arc-length along the section: this equation corresponds to the first term of the Neumann series, and is the exact solution for this reflex. This suggests that, at least for certain simple geometries, this reflex can be a source of shape information in a self shadowing region, where shape from shading techniques are not applicable. We can employ the result of section 2.4.2 to see that, referring to the geometry and terms in figure 21, this equation is:

$$N(x) = N_0 \frac{\rho}{2} \left[\frac{s + s^* \cos(\alpha)}{(s^2 + s.s^*.\cos(\alpha) + s^{*2})^{\frac{1}{2}}} \right]_{s^*=L_{b0}}^{s^*=L_{b1}} \quad (4)$$

[1] This was unnecessary when we were considering a corner, because the initial radiance on each plane due to a point source at infinity is constant, and isotropic ambient illumination simply adds an offset.

where $\alpha(s) = \frac{\pi}{2} - \frac{s}{r}$, $s = r.cos^{-1}(1 - \frac{x}{r})$ and $0 \leq x \leq r$.

This equation accurately predicts the shape of the reflexes observed (see figure 24 and figure 26), and places strong constraints on the shape of a surface within a self shadowed region, in which conventional shape from shading techniques cannot operate. We believe, therefore, that cues of this type may be extremely valuable to programs with advanced shape perception skills. In particular, Koenderink and Van Doorn [8] note that artists are accustomed to reproduce this reflex in an exaggerated form as a shape cue, which suggests that humans may employ it in some way.

Since the reflex occurs only in the self shadow region of the cylinder, we may also raise the intriguing possibility that one may use such effects to discover self shadow regions.

4 Implications for machine vision

We have shown the typical effects of mutual illumination in real images. In particular, it is easy to see that mutual illumination effects provide a strong cue for absolute surface lightness estimation, and can confound shape perception if it is based solely on reflectance map techniques.

Our data seem to indicate that the most important effects occur in concave regions. So far, we have succesfully employed a simple lighting model of the combination of a source at infinity and ambient illumination to obtain our results. From observations of the radiance distributions in cast shadow regions of our images, we feel that isotropic ambient illumination is probably a fairly good model of global mutual illumination effects, as the radiance in these regions appears to be constant, and relatively independant of surface orientation.

In this paper we have demonstrated mutual illumination effects at the intersection of two planes, at the intersection of a plane and a convex cylinder, in a smooth concave region, and in a complex polyhedral scene. This demonstrates some effects that can result from mutual illumination in different geometries:

Isotropic ambient illumination

This effect, which we showed in the blocks world images, appears to be typical of complex scenes containing only light objects. Mutual illumination appears to be an important source of isotropic ambient illumination.

Step edges with a reflex

These signatures occur at concave polyhedral edges. The size of the effect provides a cue to lightness. These edges are an attractive example of the interaction of shape and lightness cues. For obvious reasons, these reflexes do not occur at convex edges. Clearly, edge detectors that are optimised to respond to step edges may not detect certain object boundaries on such scenes, and will miss much important information. It is important, therefore, to concentrate effort on building edge detection systems that will respond to a range of features, in the vein of the ideas described in, for example, [2].

Reflexes in self shadow regions

Strong reflexes occured in the self shadow regions of a convex cylinder on a white planar background. We have demonstrated that these reflexes contain potentially strong shape cues.

It should be clear that the dependence of mutual illumination effects on scene geometry is such that one cannot use a reflectance map technique to account for them. A good example of this failure is that of the intersection of two planes, where patches with a constant surface normal have different radiance values. This implies that the photometric invariants described by Koenderink and Van Doorn [7] do not apply to objects where mutual illumination effects manifest themselves, as this approach relates the singular behaviour of the field of isophotes to that of the Gauss map by assuming that radiance is purely a function of the Gauss map. Furthermore, the fact that radiance at a point is not purely a function of the surface normal at that point implies that conventional photometric stereo techniques (Horn's book [6] is a good general reference for these techniques) will produce inaccurate shape representations for many objects.

It appears to us that the most productive approach to solving the considerable problem of recovering shape information from radiance effects involves constructing a "dictionary" of the most common generic mutual illumination effects, and using this dictionary to recognise and estimate surface shape. This work leaves a number of issues open; amongst the most important are:

- To what extent is the first order term in the Neumann series a useful approximation to mutual illumination effects, that is, to the whole series?
- We have demonstrated a geometry (convex cylinder on a plane background) where a reflex is essentially a first order effect. Is this true of many geometries? Can geometric information be extracted from the shape of this reflex?
- Is it possible, if only under restricted circumstances, to recover surface shape using the mutual illumination equation?
- What constraints on surface lightness and on shape can be recovered using mutual illumination in less restricted geometries than those shown?
- To what extent does the human visual system recognise and exploit information derived from mutual illumination?

5 Conclusion

We have derived the equation governing mutual illumination, demonstrated numerical solutions to mutual illumination problems, and have shown that these solutions are similar to effects occurring in real pictures. Our data imply that shape from shading based on the image irradiance equation may make real errors on images of concave objects. Furthermore, our data indicates that edge detectors that respond to only step edges will perform badly on polyhedral scenes, and will waste information.

We have discussed the cues that mutual illumination effects provide the vision system, and shown how its effects may help achieving decisions about shape and about surface lightness.

6 Acknowledgements

We are most grateful to Andrew Blake for a number of helpful discussions and insightful suggestions.

We thank Michael Brady, Chris Brown and Joe Mundy for help, discussions and encouragement. Paul Hubel provided measurements of relative surface lightness.

References

[1] Brady, J.M. and Ponce, J. "Toward a Surface Primal Sketch," MIT AI-Memo 824, MIT AI Lab, 1985.

[2] Canny, J.F. "Finding edges and lines in images," MIT TR-720, MIT AI Lab, 1983.

[3] Cohen, M.F. and Greenberg, D.P. 'The Hemi-cube: A Radiosity Solution for Complex Environments," *SIGGRAPH '85*, **19**, 31-40, 1985.

[4] Gilchrist, A. L., "The Perception of Surface Blacks and Whites," *Scientific American*, bf 240, 112-124, 1979.

[5] Horn, B.K.P. "Understanding image intensities," *Artificial Intelligence*, **8**, 201-231, 1977.

[6] Horn, B.K.P. *Robot Vision*, MIT press, 1986.

[7] Koenderink, J.J. and Van Doorn, A.J. "Photometric invariants related to solid shape," *Optica Acta*, **27**, 981-996, 1980.

[8] Koenderink, J.J. and Van Doorn, A.J. "Geometrical modes as a general method to treat diffuse interreflections in radiometry," *J. Opt. Soc. Amer.*, **73**, 843-850, 1983.

[9] Tricomi, F.G. *Integral Equations*, Dover, 1985.

Figures

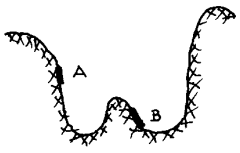

1. For this object, there is no order ρ contribution from facet A in the Neumann series for the radiance at facet B. However, there is a contribution of order ρ^2.

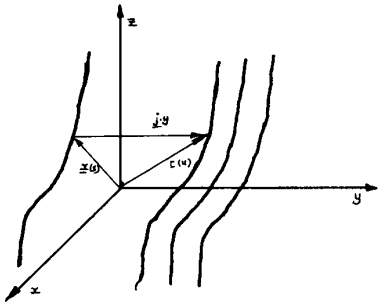

2. Terminology and geometry for the derivation of the form of the kernel for a single translational symmetry along y.

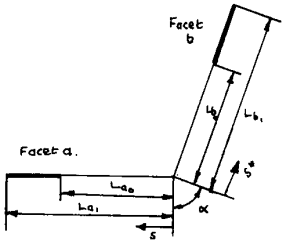

3. Geometry and terms for section 2.4.2, showing two planar facets at an angle α.

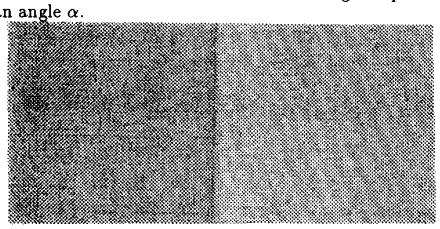

4. An image of a white 90 degree corner. The illuminant was directed off the angle bisector. The effects of the reproduction process make it difficult to see the haloes which are normally perceived on either side of the corner.

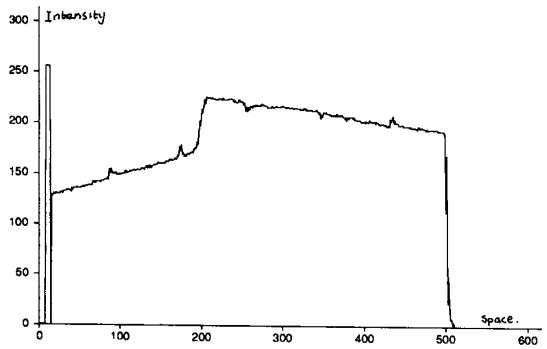

5. Image intensity observed across a white 90 degree corner, for an illuminant directed off the angle bisector. This is the typical shape for a concave polyhedral edge.

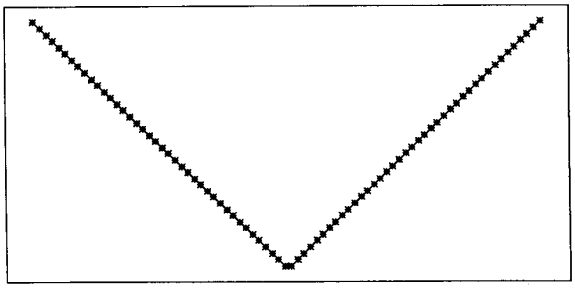

6. Geometry and nodes used to predict image intensity for a 90 degree corner.

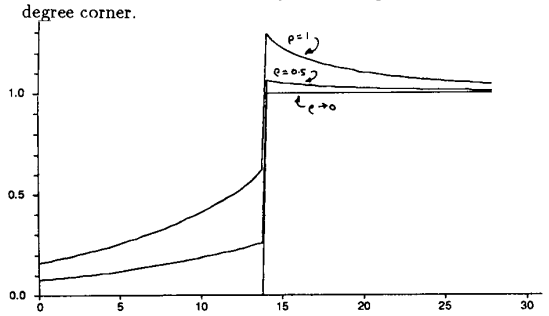

7. Image intensity predicted by the finite element method for a white 90 degree corner with the illuminant directed parallel to one plane, for a range of albedoes. Note that the reflexes become more pronounced with albedo.

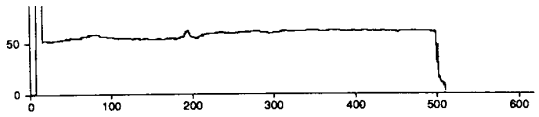

8. A specimen control: image intensity observed across a black 90 degree corner, for an illuminant directed along the angle bisector. The control images ensure that effects observed were not simply a product of light source effects. These images involved imaging black corners, identical to the white corners imaged, using exactly the same imaging geometry and light source as used in imaging the white corner. All the control images of corners have this form, an essentially flat radiance profile. The small bump in the black corner image occurs because the corners were constructed by folding cardboard, with the result that there is a line of surface facets perpendicular to the camera.

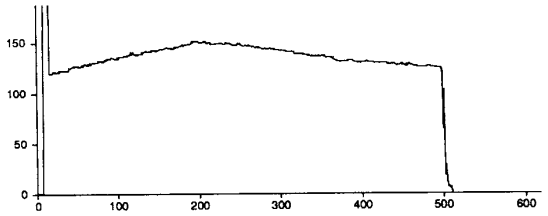

9. Image intensity observed across a white 90 degree corner, similar to figure 4, but with the illuminant directed along the angle bisector. Note the pronounced roof edge.

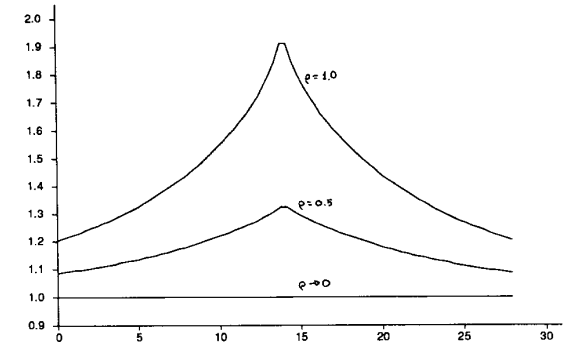

10. Image intensity predicted using the finite element method for a white 90 degree corner with the illuminant directed along the angle bisector, for a range of albedoes. Note that the reflexes become more pronounced with albedo. The flattened top to the reflexes shown is an artifact of the display technique resulting from linking the nodal values of the constant nodes employed in the finite element program.

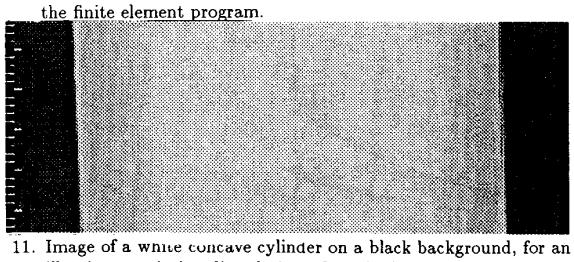

11. Image of a white concave cylinder on a black background, for an illuminant pointing directly into the cylinder.

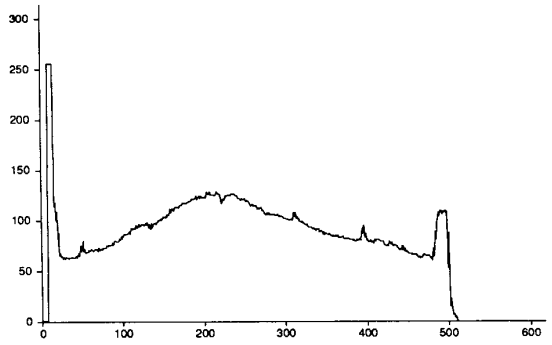

12. Observed radiance profile for a black concave cylinder on a black background.

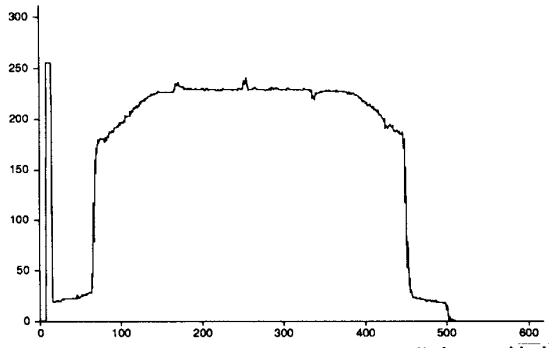

13. Observed radiance profile for a white concave cylinder on a black background. Note the pronounced flattening of the radiance profile compared with that in figure 12.

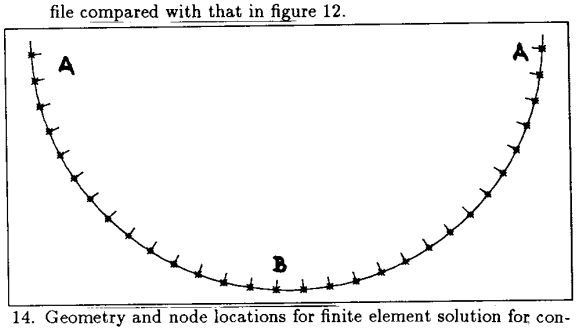

14. Geometry and node locations for finite element solution for concave cylinder.

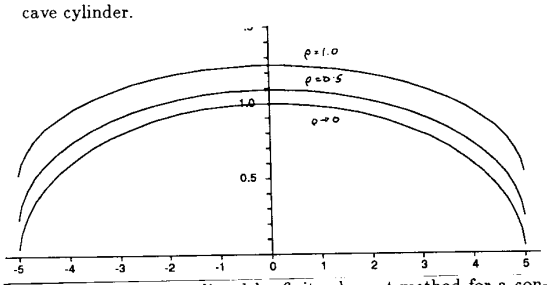

15. Radiance profiles predicted by finite element method for a concave cylinder illuminated by a point source at infinity, directed along the normal to the deepest part of the cylinder, for a range of albedoes. Note the pronounced flattening with increasing albedo of the radiance profile.

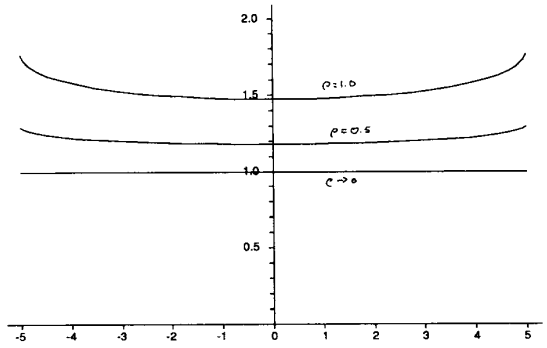

16. Radiance profiles predicted by finite element method for a concave cylinder under isotropic ambient illumination, for a range of albedoes.

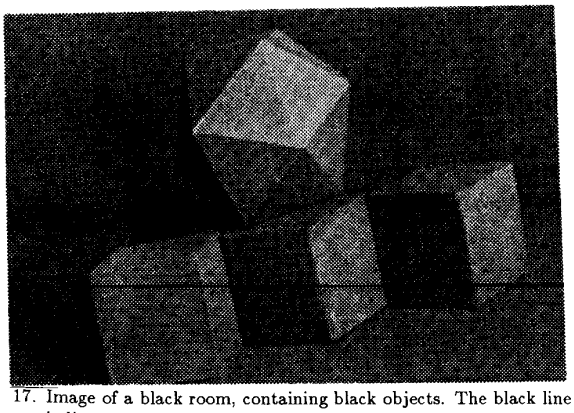

17. Image of a black room, containing black objects. The black line indicates the section represented by figure 19.

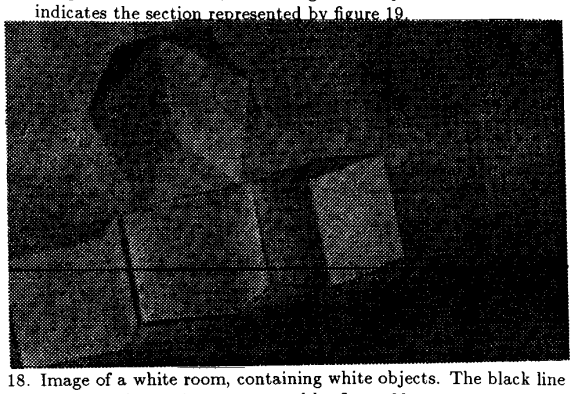

18. Image of a white room, containing white objects. The black line indicates the section represented by figure 20.

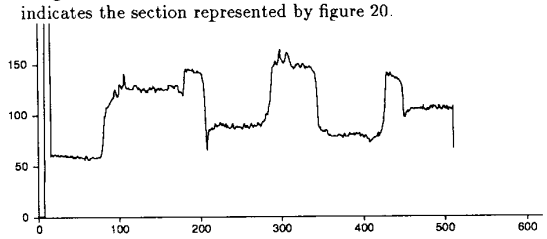

19. Section of image intensity for figure 17.

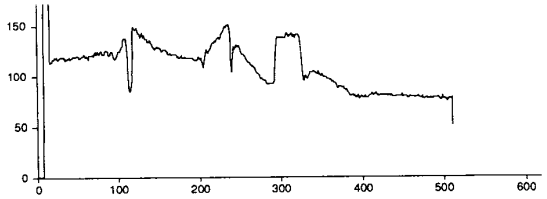

20. Section of image intensity for figure 18. Notice just how pronounced the effects of the reflexes are.

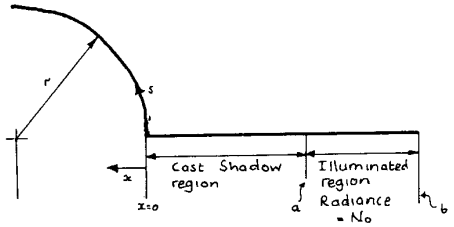

21. Geometry and terms for equation 4.

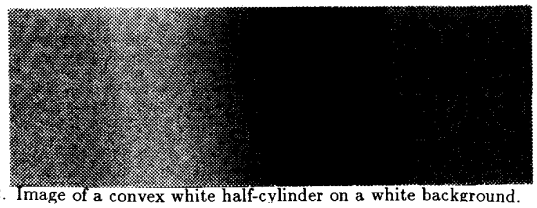

22. Image of a convex white half-cylinder on a white background.

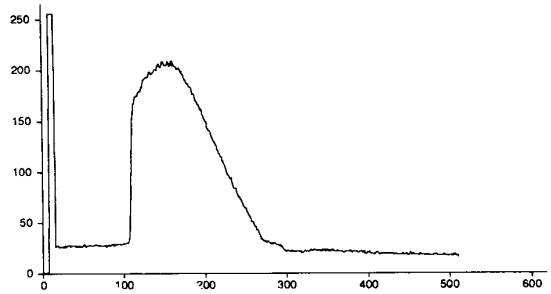

23. Image intensity measured for a convex white half-cylinder on a black background.

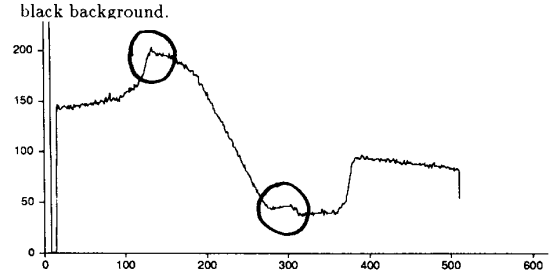

24. Image intensity measured for a convex white half-cylinder on a white background. Note the significant reflexes at both cylinder boundaries, which are circled.

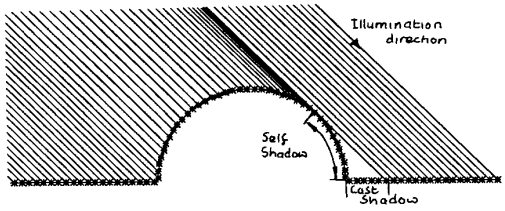

25. Geometry and node locations for finite element method for convex half-cylinder on planar background, showing illuminant direction, self shadowed region, and cast shadow.

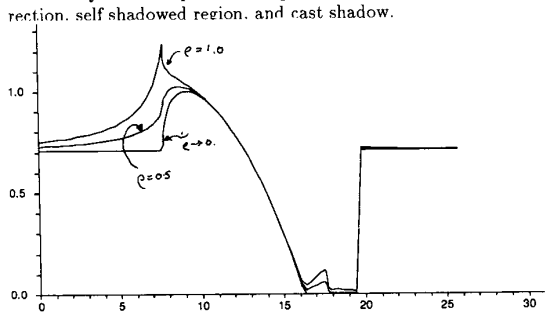

26. Image intensity predicted using finite element method for a convex white half-cylinder on a white background, for a range of albedoes. There is very good agreement with the intensity observed for this geometry (figure 24).