Shape from shading in the light of mutual illumination

David Forsyth and Andrew Zisserman

Oxford University, Oxford.

Shape from shading schemes are based on the assumption that image radiance is a function of surface normal alone. Unfortunately, because surfaces illuminate one another, radiance is a complicated global property of surface shape.

We review briefly the equations governing mutual illumination effects, and demonstrate that mutual illumination forms a major component of image radiance. We then discuss the consequences of mutual illumination effects for different theories of recovering shape from radiance. We show that discontinuities in image radiance originate solely in surface discontinuities, from shadows and from changes in surface reflectance. We argue that for a large class of shapes the response of edge detectors will also remain unchanged. Our proof involves a bookkeeping method that applies with minor modifications to discontinuities in derivatives of the radiance.

We argue that edges must be an important shape cue, because they bear a tractable relationship to three dimensional shape. The radiance of a surface cannot be of such importance, because it is intractably coupled to shape. We suggest that, because discontinuities in the derivatives of radiance arise in a highly structured fashion, they may contain exploitable shape cues.

It has been argued for many years that human subjects recover three dimensional shape cues from the distribution of shading on an object. Many machine vision researchers have interpreted this skill as an indication that the distribution of radiance in regions of an object that do not lie in shadow, can be integrated to yield a dense depth map. Unfortunately, this belief is critically dependent on an extremely simple photometric model, known as the image irradiance equation (see [11] for a clear exposition of this approach).

This model is flawed because it assumes that radiance is a function of a purely local geometric property, the surface normal. It ignores the fact that patches of surface reflect light not only to an imaging sensor, but also to other patches of surface (an effect known as "mutual illumination"), making the distribution of radiance a complicated function of the global scene geometry.

Unfortunately, it is easy to observe the effects of this redistribution in simple experiments [18,4,16,12,13]. We show some examples in figures 1-3. At first glance, the effects of mutual illumination appear rich in desirable cues to shape and to absolute surface lightness. However, the mathematical complexity of mutual illumination effects means that it is very difficult to see how these cues may be exploited. On the other hand, there is reason to believe that humans can exploit these cues to some extent [6,7].

The purpose of this article is to explore the implications of mutual illumination and to consider what shading cues can in fact be used to recover shape information. We show that although radiance itself is not a reliable shape cue, *discontinuities* in radiance are, as they can appear only at distinguished points on surfaces. Furthermore, the events that cause discontinuities in radiance are geometrically simple.

1 Mutual illumination physics

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 $\mathbf{u}$, denote the radiance by $N(\mathbf{u})$. Write $\rho(\mathbf{u})$ for the albedo of the surface patch parametrised by $\mathbf{u}$.

Assume that all surfaces are Lambertian for the present. Although our result on continuity requires only that the bidirectional reflectance distribution function of all surfaces be smooth, the Lambertian assumption simplifies the analysis considerably. We will explore the implications of relaxing this assumption in future papers.

The radiance at a point is the sum of two terms: the radiance resulting from the illuminant alone, and the radiance resulting from light reflected off other surface patches. The second term must be a sum over all surface patches. Thus, the radiance at $\mathbf{u}$ can be written:

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

where D refers to the whole domain of the parametrisation, $K(\mathbf{u}, \mathbf{v})$ represents the geometrical gain factor (often called a form factor) for the component of radiance at $\mathbf{u}$ due to that at $\mathbf{v}$, and $N_0(\mathbf{u})$ is the component of radiance at $\mathbf{u}$ due to the effects of the source alone. Conservation of energy requires K to be symmetric. In what follows we use the term "initial radiance" to refer to N_0 . Equation 1, which expresses energy balance, is called the radiosity equation (see, for example, [16] or [3]).

When $\rho(\mathbf{u})$ is constant, equation 1 is a Fredholm equation of the second kind, and K is referred to as the kernel of this equation.

For Lambertian surfaces where only diffuse reflection occurs, the kernel takes the form:

$$K^*(\mathbf{u}, \mathbf{v}) = \frac{1}{\pi} \frac{\langle \mathbf{n}(\mathbf{v}), \mathbf{d}_{uv} \rangle \langle \mathbf{n}(\mathbf{u}), \mathbf{d}_{vu} \rangle}{\langle \mathbf{d}_{uv}, \mathbf{d}_{uv} \rangle^2} \text{View}(\mathbf{u}, \mathbf{v})$$

where $\mathbf{n}(\mathbf{u})$ is the surface normal at the point parametrised by $\mathbf{u}$, $\mathbf{d}_{uv}$ is the vector from the point parametrised by $\mathbf{u}$ to that parametrised by $\mathbf{v}$, and $\text{View}(\mathbf{u}, \mathbf{v})$ is 1 if there is a line of sight from $\mathbf{u}$ to $\mathbf{v}$, 0 otherwise and $\text{View}(\mathbf{u}, \mathbf{u}) = 0$. View is clearly symmetric, and discontinuous.

We have demonstrated elsewhere [4] that solutions of this equation are in good qualitative agreement with observed mutual illumination effects. This equation is the model of diffuse radiance formation in Lambertian surfaces that we use from now on, and is the basis for the propositions proved below.

There are several approaches for solving equations of this type (see, for example, [19]). One well-known solution for the case of constant reflectance ρ is given by a Neumann series:

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

Where

$$K_n(\mathbf{u}, \mathbf{v}) = \int_D K(\mathbf{u}, \mathbf{w}) K_{n-1}(\mathbf{w}, \mathbf{v}) d\mathbf{w} \text{ and } K_1 = K$$

It is easy to prove that this geometric series always converges for $\rho < 1$. The n 'th term in this series corresponds to the contribution to radiance of a ray that has been reflected $n+1$ times, and evaluating a partial sum of this series corresponds to a form of ray tracing which admits diffuse reflection. In general, although the rate of convergence of the series depends on ρ and on the particular shape, the first few terms of the series give insight into the extent of mutual illumination effects. For example, mutual illumination effects will prove significant in concave patches. Notice also that the operator that takes the initial radiance to radiance is *linear*, so that the effect of superposing a set of sources is found by adding their individual effects.

2 Traditional Shape Cues

Several different features in image radiance have been employed to recover shape properties of surfaces. Most of these techniques require radiance to be local, and fail in the presence of mutual illumination. We discuss the effects of mutual illumination on each of the commonly used features.

2.1 The radiance itself

Conventional shape from shading ([11], and many others) explicitly computes the surface normal at each point of a shape by solving a partial differential equation (the "image irradiance equation") in the normal and the radiance. This approach clearly requires that radiance be a function of surface normal alone.

Figures 1 and 2 confirm that mutual illumination causes significant qualitative changes in the radiance. As a result of the global nature of mutual illumination effects, to construct a quantitative reconstruction of a shape with known reflectance map from its radiance under a known source, it is necessary to account for the effects of far off patches of the shape in estimating the surface normal of any patch. Under certain circumstances it may be possible to ignore distant patches or approximate their effects by isotropic ambient illumination.

If the solid angle subtended at a point by a distant surface is not small, it may make a large contribution to radiance at that point. Thus, for example, a large white wall may make significant contributions to the radiance of a scene, although not visible to the imaging device. Even under controlled conditions, unless all surfaces are nearly black or consist of an isolated convex surface, the effects of mutual illumination are significant.

As a result, quantitative shape from shading is intractable in any but the simplest cases: to succeed at quantitative shape from shading, we may have to account for things we cannot even see. Furthermore, even if the sensor can see all radiating surfaces, their interaction is complex, global and non-linear.

2.2 Singularities in radiance

If radiance is given by a function of the surface normal alone, then singularities of the Gauss map [17] will cause specific events to occur in the field of isophotes. For example, a saddle point in radiance can occur only on a parabolic line

on the surface. Koenderink and Van Doorn's paper [15] is the best known exposition of this approach.

It is clear that mutual illumination means that the singularities in the Gauss map are no longer tightly coupled to radiance. Techniques that use this coupling to infer shape information are therefore likely to be misleading if applied to real objects.

2.3 Multiple views

Photometric stereo [11] is a technique where, by observing two images of an object, each illuminated by a source an infinite distance away in different directions, one recovers the surface normal at every visible point on the object. This technique critically depends on the assumption that the radiance at a point is a function only of the surface normal at that point. As a result, mutual illumination effects will cause photometric stereo programs to estimate surface normals incorrectly. Another effect (see [12,13], for example) is that gain due to mutual illumination can cause some surface patches to be brighter than is consistent with the photometric model. It is not clear to what extent such errors in fact affect the usefulness of photometric stereo for tasks such as model matching or grasping, which can be relatively robust. Any attempt to use the dense surface normal maps recovered by photometric stereo without reference to models will have to confront these sources of error, however.

2.4 Specularities

The movement and monocular shape of specularities provide strong cues to and constraints on, local surface geometry [1,2,10,20]. Our result, given below, that mutual illumination does not give rise to discontinuities is valid only for diffuse reflection as specularities can mimic point sources and cast shadows. In more perverse cases, multiple specular reflections can cause apparent specularities which will confuse processes, such as those of [1], that extract shape from the motion of a specularity and an assumed light source.

3 Reliable shape cues

For the case of diffuse reflection and piecewise smooth surfaces and albedo we have proved the following results (the details of the proof appear in the appendix):

3.1 Discontinuities in radiance

Mutual illumination generates discontinuities only at points where the first derivative of the surface is discontinuous or at discontinuities in surface reflectance. Discontinuities in the initial radiance are preserved in the sense that their position is unaltered (though their magnitude may well be).

Thus new discontinuities can be created by mutual illumination (see figure 3 for an example of this), but these discontinuities *cling to surface features*. Discontinuities in the initial radiance may be local (for example, changes in albedo) or global (for example, shadows) in origin. Mutual illumination can accentuate, diminish, or even, for highly unlikely geometries, null discontinuities which are local in origin.

These are strong results, because these events are easy to interpret. Thus, although one can draw few or no conclusions from observations of radiance, because of its global origins,

discontinuities in radiance offer robust, reliable and tractable shape cues.

These results arise from theoretical considerations. In practice the detection of discontinuities is carried out by edge detectors and there might be some change in the position of the discontinuity due to the mutual illumination effects.

3.2 Discontinuities in radiance derivatives

Mutual Illumination can create new discontinuities in the derivative. Discontinuities in the derivative of the initial radiance are in general preserved.

Discontinuities in one of the derivatives of initial radiance occur, for example, at self shadows - where the illuminant direction grazes the surface. The new discontinuities in the derivative of radiance created by mutual illumination are essentially occlusion effects (see figure 4 and the appendix). In general these phantom self shadows will not null the discontinuities in the initial radiance derivative.

The method of proof of the first result allows us to show that discontinuities in the derivatives of radiance that arise from mutual illumination appear in a highly ordered fashion (see the appendix). In particular, they occur in a predictable way when surfaces occlude one another. This encourages speculation that these discontinuities could also represent shape cues, although it is difficult at this stage to say how they might be employed. Informal experiments indicate that these discontinuities are often small and hard to localize.

4 Conclusion

Mutual illumination is an important source of radiance that confounds simple attempts to extract shape features from a radiance signal. The complexity of the relationship between radiance and shape makes quantitative shape from shading appear impossible. However, certain features, such as step edges and self shadow boundaries, are very largely immune to the effects of mutual illumination, and as a result are of much greater importance to real vision than exact measurements of absolute radiance. Attempts have been made to exploit shadow cues. These techniques, however, approximate a dense depth map from the shadows obtained under many different illuminants. See, for example, [8,9,14].

However, to exploit these features properly, we need to change the way in which we think about shape representation. In particular, trustworthy dense depth or surface normal maps cannot be obtained using traditional techniques. Shadow and edge information is sparse and qualitative in nature, but reliable.

A great deal of work is necessary to determine how these cues may be extracted from images and how they may be exploited.

Appendix

In this section we present a sketch of the proof that discontinuities in radiance occur only at surface features or shadows, and show how to account for discontinuities in the derivatives of radiance. We have omitted unifying mathematical details here, but have given sufficient detail that it is possible to reconstruct a full proof. The key idea is to use the radiosity equation (equation 1) to chart the discontinuities in radiance and its derivatives. This is accomplished by

comparing the radiance and its derivatives at neighbouring points on the surface.

For simplicity we consider an arrangement of surfaces and illuminants which has translational symmetry. Consequently it is only necessary to deal with the surface profile. The graphical aid presented below in the proof captures all the essential geometric features, and the extension to more general surfaces is straightforward.

Translational symmetry in the arrangement of surfaces and illuminants makes it possible to integrate out one degree of freedom in the kernel, and the radiosity equation (1) reduces to:

$$N(s) = N_0(s) + \rho(s) \int_{\Gamma} k(s, s') View(s, s') N(s') ds' \quad (2)$$

where Γ is the cross section and s is arc length. The term $k(s, s') View(s, s')$ is the original kernel with the contribution along the symmetry integrated out. We use the notation $N = N_0 + \rho KN$ for equation 2.

Discontinuities in radiance

$N(s)$ is discontinuous at s if

$$\Delta N(s) = \lim_{\delta \rightarrow 0} |N(s + \delta) - N(s - \delta)| \neq 0$$

ΔN provides a measure of the magnitude of the discontinuity. For the proof we assume that the radiance obeys equation 2 and is bounded.

We first examine the discontinuities in KN . These arise because the kernel $k(s, s')$ is discontinuous. There are 2 types of discontinuity - creases (where $k(s, s')$ is discontinuous) and changes in occlusion (where $View(s, s')$ switches between 0 and 1). Although KN is necessarily continuous when $k(s, s') View(s, s')$ is, discontinuities in the kernel do not necessarily generate discontinuities in KN . Analysing where discontinuities occur is simply visualised by examining a representation of the discontinuities in the kernel, called the Crease Occlusion Graph (COG). A typical COG is shown in figure 5.

KN is discontinuous if ΔKN is not zero. This can only occur when the domain over which the integrand (i.e. $k(s, s') View(s, s') N(s')$, for fixed s) is non-zero (the support of the integrand) changes sharply. This domain is the pool of ones in $View(s, s')$, for fixed s . In figure 6, the horizontal lines represent the domain of integration (with respect to s' for fixed s), and the intersections between the lines and the shaded regions represent the support of the integrand.

KN can be discontinuous if s is at a crease because here the support of the integrand changes sharply, causing a sharp change in KN (see line c in figure 6). For most points, KN is continuous even though the kernel is not, because the support of the integrand changes only slightly for a small change in s (line a in figure 6). In particular, notice that when a change in s causes a new pool of ones to appear in the support of the integrand (line b figure 6), a discontinuity can only be generated when the pool itself grows discontinuously, i.e. changes suddenly from the empty set to an open set¹, which happens when the boundary of a pool of ones is locally parallel to the integrating line. This can happen with piecewise smooth surfaces only when the surface has a planar

¹That is, to a set of non-zero measure. The core of this argument is that continuity of KN is guaranteed by a measure condition. For a detailed presentation, see [5].

patch. However, this does *not* generate a discontinuity because along this patch the kernel itself falls to zero over the parallel section of boundary, because this section represents the case when the planar patch is viewed in its plane. Thus, KN is discontinuous only at creases.

We have seen that KN is continuous except at creases. Thus, ρKN is in general discontinuous at creases, and at discontinuities in ρ . Special cases where a discontinuity in ρ lies on a crease, and the discontinuity in ρ nulls that at the crease can be disturbed by a small change in either ρ or the geometry, and are therefore highly unlikely. Thus, for N to be discontinuous, we must have one of the following cases:

1. N_0 discontinuous, ρKN continuous

This case occurs at surface points like point C in figure 3 where KN is continuous but there is a discontinuity in the initial radiance due to, for example, a cast shadow. Here because its effects are continuous, mutual illumination does not affect the magnitude of the discontinuity (though the magnitude of the radiance itself may well be altered).

2. N_0 continuous, ρKN discontinuous

This can occur, for example at points like B in figure 3, where the surface is in shadow, but would generate a discontinuity were it illuminated. Here N is discontinuous as a result of mutual illumination, *but the discontinuity must lie on a crease or a discontinuity in albedo*.

3. Both discontinuous

This is the general case for creases or discontinuities in albedo. This proof is concerned with the *existence* of discontinuities, but has no bearing on their magnitudes or signs, so that we cannot say that either sign or magnitude will be preserved. It is unlikely, however, that discontinuities will have the same magnitude and exactly cancel.

In summary:

Discontinuities are introduced by mutual illumination only at creases and discontinuities in albedo, which are surface features. Discontinuities in the initial radiance away from surface features, such as those caused by shadow, are always preserved. Discontinuities in the initial radiance at creases and at discontinuities in albedo will be preserved in general - but their magnitude (not position) may change.

The rest of the proof consists of checking that these results hold for more general surfaces. The details appear in [5].

Discontinuities in radiance derivative

The argument for preservation/creation of discontinuities in radiance derivative is similar to the above. In this case, we assume that the geometry and source are such that radiance and initial radiance are piecewise differentiable. However, instead of using equation 2 for the radiance, we consider its derivative, $\frac{d}{ds}N(s)$.

$$\frac{dN(s)}{ds} = \frac{dN_0(s)}{ds} + \frac{d\rho(s)}{ds}KN + \rho \frac{d}{ds}\{KN\} \quad (3)$$

We must admit generalised functions for the derivatives to be meaningful. We assume that the sum of two discontinuous functions is discontinuous; this is true for almost all

functions. In particular, it will be true for almost all geometries. Hence, the derivative of radiance will be discontinuous when any single term is discontinuous. Inspection of the first two terms in equation 3 and a repeat of the process described above shows immediately that the derivative of radiance will be discontinuous when the derivative of the initial radiance is discontinuous and when the derivative of the albedo is discontinuous ($\frac{d}{ds}N(s)$ is undefined at discontinuities in albedo and at creases where N is discontinuous).

The COG can be used to discover discontinuities in the fourth term.

$$\rho \frac{d}{ds}\{KN\} = \rho(s) \int_D \frac{\partial\{k(s,s')View(s,s')\}}{\partial s} N(s') ds'$$

In general $\frac{\partial}{\partial s}\{k(s,s')View(s,s')\}$ consists of delta functions along the borders of the pools of ones in the graph, except where the border is parallel to the s direction. Elsewhere, except at creases, it is continuous. Discontinuities in this term can occur in two possible ways; a horizontal line touches the boundary of a pool in the COG, or a discontinuity in N is "sampled" by one of the delta functions.

The first case depends on the geometry of the pools in the COG; an example is shown in figure 7 (line d). The second case is easily predicted from the COG, because, as we have seen, we can predict the discontinuities in N (see figure 7, line e). This method can clearly be applied to higher derivatives of the radiance.

Acknowledgements

We thank Andrew Blake for numerous stimulating discussions. We thank Michael Brady for his support and interest. Chris Brown, Margaret Fleck, Joe Mundy, and Guy Scott have made many helpful suggestions. DAF thanks Magdalen college for financial support. AZ thanks the SERC for financial support.

References

- [1] Blake, A. and G. Brelstaff "Geometry from specularly," *Proc 2nd ICCV*, 394-403, 1988.
- [2] Brelstaff, G. and A. Blake "Detecting specular reflections using Lambertian constraints," *Proc. 2nd ICCV*, 297-302, 1988.
- [3] Cohen, M.F. and Greenberg, D.P. "The Hemi-cube: A Radiosity Solution for Complex Environments," *SIGGRAPH '85*, 19, 31-40, 1985.
- [4] Forsyth, D.A. and Zisserman, A. "Mutual Illumination," *Proc. CVPR*, 1989.
- [5] Forsyth, D.A. and Zisserman, A. "Shadows and light: mutual illumination, discontinuities and shape," Oxford University Department of Engineering Science Internal Report, *In preparation*, 1989.
- [6] Gilchrist, A. L., "The Perception of Surface Blacks and Whites," *Scientific American*, bf 240, 112-124, 1979.
- [7] Gilchrist, A. L. and A. Jacobsen "Perception of lightness and illumination in a world of one reflectance," *Perception*, 13, 5-19, 1984.
- [8] Hatzitheodorou, M.G. and Kender, J. "An optimal algorithm for the derivation of shape from darkness," *Proc. CVPR*, 1988.

- [9] Hatzitheodorou, M.G. "The derivation of 3D surface shape from shadows," Proc. DARPA Image Understanding, 1989.
- [10] Healy, G. "Local shape from specularity," *Proc. 1st ICCV*, 151-160, 1987.
- [11] Horn, B.K.P. *Robot Vision*, MIT press, 1986.
- [12] Ikeuchi, K., H.K. Nishihara, B.K.P. Horn, P. Sobalvarro and S. Nagata "Determining grasp configurations using photometric stereo and the PRISM binocular stereo system," *IJRR* 5, 1, 1986.
- [13] Ikeuchi, K. "Determining a depth map using a dual photometric stereo," *IJRR* 6, 1, 1987.
- [14] Kender, J. and Smith, E.M. "Shape from darkness; deriving surface information from dynamic shadows," Proc. AAAI 1986.
- [15] Koenderink, J.J. and Van Doorn, A.J. "Photometric invariants related to solid shape," *Optica Acta*, 27, 981-996, 1980.
- [16] 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.
- [17] O'Neill, B. *Elementary Differential Geometry*, Academic, 1966.
- [18] Pearson, D.E and Robinson, J.A. "Visual communication at very low data rates," *Proc. IEEE*, 74, 4, 795-812, 1985.
- [19] Tricomi, F.G. *Integral Equations*, Dover, 1985.
- [20] Zisserman, A., P. Giblin and A. Blake "The information available to a moving observer from specularities," *Image and Vision Computing*, 7, 1, 38-42, 1989

Figure 1: When ρ is small, solutions of the mutual illumination equation tend towards the radiance predicted by the image irradiance equation as mutual illumination effects will be dominated by source effects. Thus the effects of mutual illumination can be observed experimentally by comparing the radiance of a white set of objects and a black set. Qualitative differences in radiance distributions for images taken of similar arrangements of these objects can be ascribed to the effects of mutual illumination. This figure shows a section of the radiance observed for a black cylindrical gutter cut in a black plane, illuminated from above. Radiance features marked A and B in this and the following figure occur at the points on the profile (shown in the inset) marked with corresponding letters.

Figure 2: A section of the radiance observed for a white cylindrical gutter cut in a black plane, illuminated from above. Note that this signal is qualitatively very different from that of figure 1. In particular note the constant central region - this is not a saturation effect. Quantitative shape from shading, based on the image irradiance equation, would produce entirely the wrong result with this data.

Figure 3: A surface feature highlighted by mutual illumination. a) shows the geometry for this case, a white cylindrical hump on a white planar background. The surface crease labelled B is not illuminated by the primary illuminant. The shadowed surface is however illuminated by secondary reflections. b) Measured radiance for this geometry: there is a discontinuity in radiance at B, which not predicted by the image irradiance equation, and one at C, caused by a shadow. The radiance features labelled A, B and C in (b) occur at the points with the corresponding labels in (a).

Figure 4: In this simple geometry there is a discontinuity in the derivative of radiance at point A due to mutual illumination.

Figure 5: The crease occlusion graph charts the events which may cause discontinuities to appear in KN and its derivatives. We show this graph for a simple profile. A point (s, s') in the square represents a pair of points on the profile, which is drawn above and to the left of the graph. Regions where the corresponding points have a line of sight (i.e. $\text{View}(s, s') = 1$) are hatched, and regions where they do not, are left blank. The dashed lines are creases on the surface (discontinuities in surface orientation). Note the symmetry about the line $s = s'$.

Figure 6: Evaluating KN is equivalent to integrating (w.r.t. s') the radiance along a horizontal line. This is the COG of figure 5, with the shapes omitted. At $s = a$, KN is continuous, because the support of the integrand changes only very slightly for nearby lines. It is also continuous at $s = b$ - even though the variation takes the integral away from an occlusion boundary. This is because the relevant regions in the support of the integrand shrinks to a point before disappearing. At c , however, where a change in s crosses a crease in the surface, the support changes sharply, and the integral may be discontinuous.

Figure 7: If a discontinuity in N lies at $s' = f$ (dotted line), in forming $\int_{\Gamma} \frac{\partial}{\partial s} \{k(s, s') \text{View}(s, s')\} N(s') ds'$, the delta functions, which lie along the heavy border, will "sample" the discontinuity at $s' = f$, and the integral will be discontinuous at $s = e$. The integral will also be discontinuous at $s = d$, where the line of integration just touches the border, because at $s = d$ the integral contains a sample of radiance, and for $s < d$ it does not. For values of s where neither of these events occur, the integral is continuous. The COG is that of figure 5.