Shadows and Shading Flow Fields

Pierre Breton
breton@cim.mcgill.ca

Steven W. Zucker
zucker@cim.mcgill.ca

Centre for Intelligent Machines, McGill University
3480 University St., Montréal, Québec, Canada, H3X 2A7

Abstract

Many presume that parsing the shadows out of an image is a high-level task, because of the global nature of the shadow formation process. But shape-from-shading algorithms are low-level, in the sense that they seek solutions (surface normals or depth values) directly from image intensities. A dilemma arises: since shape-from-shading involves an illumination term, shadows must first be identified. We show that a structure intermediate between intensities and surfaces — the shading flow field — provides a solution to this dilemma. Our analysis is based on the observation that the geometric information that can be derived from images supports different inferences than the photometric information, and our specific goal will be to articulate this geometric structure and to show how shading flow fields can be reliably computed.

1 Introduction

Figure 1(top-left) contains an image of a sphere against a plane illuminated by a partially obscured source. Our perception of the sphere is robust against this occluded source, but classical computer vision algorithms are not (Fig. 1(top-right)). Clearly, it is not appropriate to assume that the illumination is constant across the image and similar arguments could be made for the albedo, and light source direction.

This thus raises the particular question of what does remain fixed. To illustrate our answer, in Fig. 1(bottom-left) we show the local geometric description that we shall compute for this example. It shows the shading flow field (thin arrows), in this case a direction field indicating the isophotes [8], or directions of constant intensity, superimposed on the tangent field, or local edge directions (thick arrows). These comprise what we mean by the geometric structure in the image, and they are what remain constant. Note, in particular, how smooth the shading flow field is under the cast illumination edge, and how aligned the bounding edges remain under it. That this information is sufficent to support surface computations is indicated in Fig. 1(bottom-right), although we shall not discuss that part of the computation in this short paper.

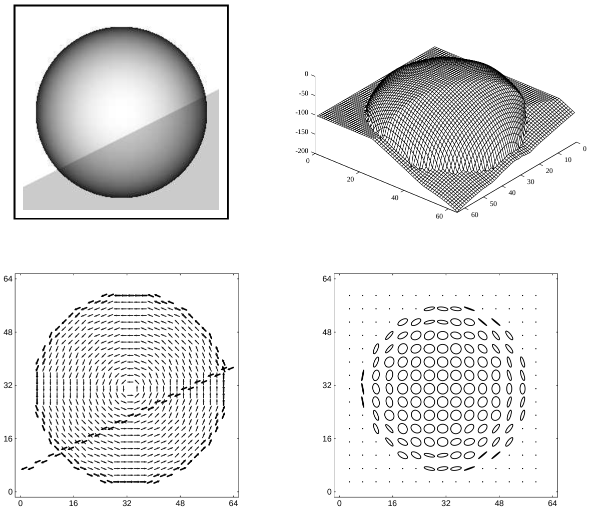

Figure 1: Top-left: an image for which assumptions of standard shape-from-shading [1] are violated. Top-right: neither of the two most plausible scenes (the surface albedo changes abruptly or a semi-transparent object lies somewhere in the scene) can be implied from the depth map obtained. Bottom-left: the shading flow field is depicted by thin arrows and the tangent field, by thick arrows. Bottom-right: the accurate shape recovery illustrated by the surface normals is obtained from the shading flow field [3].

Thus our formulation of the shape-from-shading problem generalizes the classical one, which is based on assumptions of uniform albedo and lack of shadows everywhere. As background to our analysis we stress that standard shape-from-shading algorithms (e.g. [5, 11]) are based solely on the photometric component of the image (i.e. the image intensities). From Lambert's law we obtain: $I(x,y) = \rho\lambda\mathbf{L}\cdot\mathbf{N}(x,y)$ where $I(x,y)$ is the image intensity[1] at a point (x,y); ρ is the albedo of the surface, i.e. the fraction of the shining light which is reflected; λ is the illumination, i.e. the amount of shining light; $\mathbf{L}$ is the light source direction; and $\mathbf{N}(x,y)$ is the normal at the surface point corresponding to an image point (x,y). By assuming that smooth variations in intensity are entirely due to smooth variations in surface orientation:

$$\vec{\nabla} I(x,y) = \rho\lambda\mathbf{L}\cdot\vec{\nabla}\mathbf{N}(x,y) \qquad (1)$$

the classical basis for surface reconstruction is set.

In this paper, we advocate *not* to relate the image intensity directly to surface shape as the photometric and geometric properties of the scene are then confounded, and assumptions about the relation between image intensities and image irradiance have to be made. Instead, we propose using geometric properties of the image, in particular the shading flow field, as an intermediate structure that captures image irradiance properties.

2 The Shading Flow Field

In the simplest cases, the shading flow field indicates the direction in which the image intensity remains constant. To start, consider an image domain D that derives from the projection of an unoccluded smooth surface. For $(x,y) \in D$, define $\Theta(x,y)$ by $\vec{\nabla} I(x,y) \perp \Theta(x,y)$, where $\Theta(x,y)$ is the shading flow field. We describe $\Theta(x,y)$ as carrying the geometric information precisely because it specifies a form of parallel transport and a moving Darboux frame.

The distinction between these two dual flow fields becomes apparent in the curvatures κ^s and κ^g; both are necessary, since they constrain the flow fields in orthogonal directions. The variation of the orientation of the shading flow is $\frac{\partial\theta}{\partial r} = (\kappa^s \sin\vartheta + \kappa^g \cos\vartheta)$, where the variables r and ϑ are the local frame coordinates expressed as polar coordinates. We shall derive the transport equations (through curvatures) in Sect. 3 and it is this equation that supports the reliable computation of flow fields.

[1] The image irradiance $E(x,y)$ is often replaced by the image intensity $I(x,y)$. One assumes that the relation between the image irradiance $E(x,y)$ and the image intensity $I(x,y)$ is linear. One further assumes that the highest image intensity value corresponds to a surface facing the light source so that one can use the image intensity normalized by the highest value, and take ρ and λ being given the unit value.

The shading flow field is thus locally described by $\Psi = \{x, y, \theta, \kappa^s, \kappa^g \mid (x,y) \in \mathbb{R}^2, \theta \in [0, 2\pi[, \kappa^s \in \mathbb{R}, \kappa^g \in \mathbb{R}\}$. More generally, we shall have to worry about discontinuities, which will affect the local neighbourhoods over which θ and κ_i are defined (Sect. 3).

The shading flow field captures image irradiance properties. The shading flow field is not dependent on the relation between the image irradiance $E(x,y)$ and the image intensity $I(x,y)$ to the extent that the relation is not itself dependent on position (x,y). The directions for which the image irradiance and the image intensity remain constant are the same, implying camera calibration is not an issue for determining this geometric structure.

The shading flow field is only dependent on the geometric properties of the scene. We rewrite Eq. 1: $\vec{\nabla} I(x,y) = \Phi(x,y)\,\vec{\Gamma}(x,y)$ to stress the photometric and geometric aspects of the scene, respectively:

$$\begin{aligned} \Phi(x,y) &= \rho(x,y)\,\lambda(x,y) \\ \vec{\Gamma}(x,y) &= \mathbf{L}(x,y)\cdot\nabla\mathbf{N}(x,y)\ . \end{aligned}$$

Since both the illumination and the albedo can be functions of position, the image intensity is a poor choice for initial data: the geometric information about the scene would be confounded with the photometric. Similarly for the first and second order directional derivatives of the intensity function.

The directional derivatives' information can nevertheless be combined in ways that are independent from the variable illumination and albedo. Although the following calculations are elementary, their implications are central for our model.

Theorem 2.1 *For the shape from shading problem, the orientation of the intensity gradient field (where it is defined) is independent of the variable illumination (λ) and albedo (ρ). Thus the orientation of the intensity gradient field only depends on the geometric properties of the surface $\vec{\Gamma}(x,y)$ and the lighting with respect to the viewer.*

Proof: The orientation of the gradient field θ is related to a ratio of directional derivatives:

$$\tan\theta = \frac{I_y}{I_x} = \tan\theta = \frac{\mathbf{L}\cdot\frac{\partial}{\partial y}\mathbf{N}}{\mathbf{L}\cdot\frac{\partial}{\partial x}\mathbf{N}} = \frac{\Gamma_y}{\Gamma_x}\ .$$

The orientation of the intensity gradient field is independent from the variable illumination (λ) and albedo (ρ). □

Corollary 2.1 *The shading flow field (where it is defined) is independent of the variable illumination and* *albedo, and thus only depends on the geometric properties of the surface and the lighting with respect to the viewer.*

We acquire two vector fields based on this orientation information. We can consider the normalized gradient field and the shading flow field as unit vector fields (direction fields).

Corollary 2.2 *The curvatures of both the shading flow field and the normalized gradient field are also independent of the variable illumination and albedo, and thus only depend on the geometric properties of the surface and the lighting with respect to the viewer.*

The curvature of shading flow field κ^s and the curvature of the normalized gradient field κ^g are respectively:

$$\kappa^s = \frac{2I_x I_y I_{xy} - I_x^2 I_{yy} - I_y^2 I_{xx}}{\left(I_x^2 + I_y^2\right)^{\frac{3}{2}}} ;$$
$$\kappa^g = \frac{I_x I_y \left(I_{xx} - I_{yy}\right) - I_{xy}\left(I_x^2 - I_y^2\right)}{\left(I_x^2 + I_y^2\right)^{\frac{3}{2}}} .$$

The shading flow field made the relation between geometric properties of the image irradiance and of the scene explicit. To carry the rest of the "photometric information", we define the "augmented" shading flow field. Let the augmented shading flow vector $\vec{S}$ be such that its magnitude equals the magnitude of the image intensity gradient:

$$\vec{S}(\mathbf{x}) = \left(\left|\vec{\nabla} I(\mathbf{x})\right| \cos\theta(\mathbf{x}),\ \left|\vec{\nabla} I(\mathbf{x})\right| \sin\theta(\mathbf{x})\right) .$$

The relation with the photometric properties is implicit since the magnitude of the augmented shading flow vector magnitude:

$$\left|\vec{S}(\mathbf{x})\right| = \rho\lambda \left|\mathbf{L}\cdot\vec{\nabla}\mathbf{N}(\mathbf{x})\right| = \Phi(\mathbf{x}) \left|\vec{\Gamma}(\mathbf{x})\right| ,$$

depends on both the photometric and the geometric properties of the scene.

The proof of the following proposition develops several key properties of the shading flow field, including the transport equations.

Proposition 2.1 *The shading flow field, augmented with the magnitude of the intensity gradient, can be locally described by* $\Psi^s = \{x, y, \vec{S}, \mathbf{D} \mid (x,y) \in \mathbb{R}^2, \vec{S} \in \mathbb{R}^2, \mathbf{D} \in \mathbb{R}^3\}$.

Proof: Consider a smooth shading flow field $\vec{S}(\vec{r})$. We develop a Taylor expansion of $\vec{S}(\vec{r})$ to obtain a local description and we keep only the first terms:

$$\vec{S}(\vec{r} + \vec{dr}) = \vec{S}(\vec{r}) + \frac{\partial \vec{S}(\vec{r})}{\partial \vec{r}} \vec{dr} + \ldots$$

where

$$\frac{\partial \vec{S}(\vec{r})}{\partial \vec{r}} = \frac{1}{2} \operatorname{curl} \vec{S} \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} + \frac{1}{2} \operatorname{div} \vec{S} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + \frac{1}{2} \operatorname{def} \vec{S}\, R_\varphi^{-1} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} R_\varphi$$

Since the shading flow field is related to image intensity gradient by:

$$\vec{S}(\mathbf{x}) = \begin{pmatrix} -I_y(\mathbf{x}) \\ I_x(\mathbf{x}) \end{pmatrix} = \begin{pmatrix} \left|\vec{\nabla} I(\mathbf{x})\right| \cos\theta(\mathbf{x}) \\ \left|\vec{\nabla} I(\mathbf{x})\right| \sin\theta(\mathbf{x}) \end{pmatrix} ,$$

it follows that $\operatorname{div}\vec{S} = \operatorname{curl}\vec{\nabla} I \equiv 0$, the trace of the variation of the flow field is null. Thus, three quantities suffice to describe it locally:

$$\vec{S}(\mathbf{x} + \Delta\mathbf{x}) \approx \vec{S}(\mathbf{x}) + \begin{pmatrix} -I_{xy} & -I_{yy} \\ I_{xx} & I_{xy} \end{pmatrix} \begin{pmatrix} \Delta x \\ \Delta y \end{pmatrix} . \tag{2}$$

The coefficients of this matrix relate to the three components of a DISTORTION vector as follows:

$$\mathbf{D} \equiv \begin{pmatrix} -\operatorname{def}\vec{S}\sin(2\varphi) \\ \operatorname{def}\vec{S}\cos(2\varphi) \\ \operatorname{curl}\vec{S} \end{pmatrix} = \begin{pmatrix} I_{xx} - I_{yy} \\ -2I_{xy} \\ I_{xx} + I_{yy} \end{pmatrix} .$$

Figure 2 provides a geometric picture of the relationship between the shading flow orientation and the nature of the distortion. Combined with the shading flow vector, it completes the local shading flow description space, $\Psi^s = \{x, y, \theta, \left|\vec{\nabla} I\right|, \mathbf{D}\}$. □

If we express the variation in terms of the coordinates defined by the tangent and the normal to the shading flow at $\vec{r}$, we can relate curvatures, κ^g and κ^s, and the differential invariants:

$$\frac{\partial \vec{S}(\vec{r})}{\partial \vec{r}} \vec{dr} = \begin{pmatrix} -\kappa^g & \kappa^c \\ -\kappa^s & \kappa^g \end{pmatrix} \begin{pmatrix} \Delta u \\ \Delta v \end{pmatrix} , \tag{3}$$

where $\Delta u = \Delta x \cos\theta + \Delta y \sin\theta$, $\Delta v = -\Delta x \sin\theta + \Delta y \cos\theta$, and $\kappa^c = -2I_{xy}I_xI_y - I_{xx}I_x^2 - I_{yy}I_y^2$. This third term can not be accounted for in the model without magnitude.

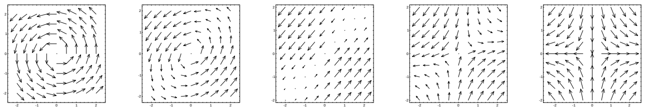

Figure 2: This sequence shows flow fields for which the distortion vector differs from a pure curl, $D_3 = 1$ (left), to a pure deformation $D_3 = 0$ (right), with the two components blending with different proportions (middle). Whereas the curl is symmetric with respect to a point, the deformation is symmetric with respect to an axis. Note how the nature of the singularities differs as the distortion changes.

3 Coherent Shading Flow Field

Since the shading flow field is the "dual" of the normalized gradient field, we begin by considering the gradient of an image. "Fuzzy derivatives" [9] cannot work since discontinuities are smoothed (Fig. 3(top-right)).

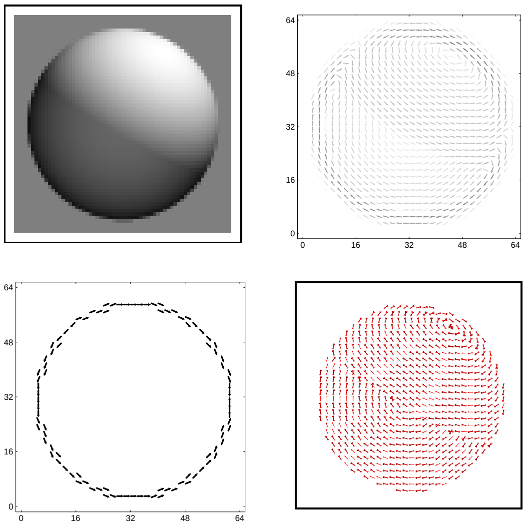

Figure 3: Top-left: an ideal intensity image of a scene illuminated by two distant point light source; top-right: the fuzzy derivatives flow field; bottom-left: the corresponding tangent field; bottom-right: shading flow field. Note the shading flow field presents again a point singularity of index one (near the centre of the upper right quadrant) and an undefined region. For display purposes only, the resolution of the edge map and shading flow field is decreased by two. 5 iterations. Again note the discontinuity in bottom-right is smoothed away in top-right.

We infer the directional derivatives of a function that captures the shading of the scene as projected onto the image plane. Key to this will be consistency properties of the shading flow field: if the underlying shading function is smooth, it should be possible to transport shading flow field descriptors along the flow to neighbouring positions. Thus our algorithm will consist of two stages: (i) initial estimation of shading flow field descriptors; and (ii) refinement of these via transport to enforce consistency.

The space of shading flow orientation is divided into a discrete set of regular intervals, Θ. Each interval defines an equivalence class which we represent with the central orientation of the interval: θ_i. Each one is viewed as a hypothesis that the directional derivative of the image intensity function is null in an orientation within the interval. Similarly, the space of possible shading flow curvatures is divided into a discrete set of regular intervals, K, which defines equivalence classes for the curvatures, κ_i^s and κ_i^g.

The shading flow field is thus redefined as: $\mathbf{\Psi} = \{x, y, \theta_i, \kappa_i^s, \kappa_i^g \mid (x, y) \in \mathbb{Z}^2, \theta_i \in \Theta, \kappa_i^s \in K, \kappa_i^g \in K\}$. The general shading flow field is a generalization of the vector fields obtained for smooth images that allows multiple values of orientation and curvature to represent discontinuities. These discontinuities will play a fundamental role in parsing shadows and albedo changes (Sect. 4).

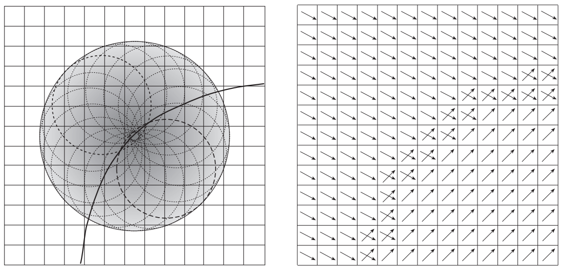

Figure 4: Representation of flow field near a discontinuity. Left: A set of 16 neighbourhoods is used to provide reliable estimates of the gradient, shown here as dotted circles. This partitioning insures that at least one fit will not be corrupted, even if there is a shading flow field discontinuity (shown here as a solid line). Note the two dashed circles provide support to the appropriate flow hypotheses in regions adjacent to the discontinuity. The confidence given a data point depends on its distance from where the model is fit, as illustrated by the radial shading. Rigth: Two distinct shading flow hypotheses should be retained to represent the flow field discontinuity.

For shading flow field discontinuities, we need to define neighbourhoods over which the shading function is smooth (see Fig. 4). Anticipating different orientations, we define a set of 16 such neighbourhoods. These neighbourhoods provide domains to fit a local model of shading function. Initial estimates are obtained by fitting a polynomial $\hat{I}(\mathbf{x}) = \sum_k a_k X_k(\mathbf{x})$ and scaling the results in proportion to a merit function. We then seek a consistent flow field out of these initial estimates.

In the discrete domain, the notion of continuity is replaced by one of consistency, i.e. local descriptors of the shading flow field at neighbouring positions (e.g. $\Psi_i^s = \{x_i, y_i, \vec{S}_i, \mathbf{D}_i\}$, $\Psi_j^s = \{x_j, y_j, \vec{S}_j, \mathbf{D}_j\}$), are samples of the same smooth shading function. To determine whether or not they are, we introduce a pairwise compatibility function for a relaxation labeling process [6] built on the transport equations defined by the flow field constraints introduced in Sect. 2.

The local descriptors Ψ_i^s of the shading flow field at a position (x_i, y_i) constrain the shading flow vector $\vec{S}_j$ at a position (x_j, y_j) through the transport equation. The differential invariants $\mathbf{D}$ (or κ^s, κ^g, and κ^c) of the augmented shading flow field capture the local variation of $\vec{S}$ (or θ, $\left|\vec{\nabla} I\right|$). The transported shading flow description $\vec{S}_i^*$ at $\mathbf{x}_j$ is derived from Eq. 2 with $\Delta\mathbf{x}$ equal to $\mathbf{x}_j - \mathbf{x}_i$ and the matrix coefficients derived directly from $\mathbf{D}_i$. Reciprocally, the local descriptors Ψ_j^s of the shading flow field at a position (x_j, y_j) constrain the shading flow vector $\vec{S}_i$ at a position (x_i, y_i), and the transported shading flow description $\vec{S}_j^*$ is obtained in a similar way.

We define the compatibility function such that it is positive only when both these constraints are satisfied. To achieve this, we use the logical/linear combinators $\curlywedge$ [7]. The compatibility is then given by $r_{ij} = r_{ij}^* \curlywedge r_{ji}^*$ where

$$r_{ij}^* = G''(\theta_j - \theta_i^*)\ G\left(\left|\vec{\nabla} I_j\right| - \left|\vec{\nabla} I_i^*\right|\right)\ G\left(\left|\vec{\mathbf{x}}_j - \vec{\mathbf{x}}_i\right|\right) \ ,$$
$$r_{ji}^* = G''(\theta_i - \theta_j^*)\ G\left(\left|\vec{\nabla} I_i\right| - \left|\vec{\nabla} I_j^*\right|\right)\ G\left(\left|\vec{\mathbf{x}}_i - \vec{\mathbf{x}}_j\right|\right) \ .$$

4 Interpretations of Discontinuities

To establish the signature of each scene discontinuity, we start with generic surfaces illuminated from a generic direction and viewed from a generic position[2]. Immediately, we have:

Observation 4.1 *A smooth matte surface of constant albedo projects, under constant illumination, a smooth image irradiance and almost everywhere, a smooth shading flow field.*

[2]The assumptions of "*generic view*" and "*generic illuminant direction*" states that the observer and the light sources are not in a special position relative to the scene. A small change in their position should not affect the qualitative aspect of the image [2, 4].

Shadows occur when some light source does not shine on part of the scene. The boundaries of shadow correspond to where the illumination abruptly changes.

Let $\lambda_{(i)}\ \mathbf{L}_{(i)}$ describe the illumination and the direction of light sources indexed by i. Consider a shadow caused by the occlusion of a point light source $\lambda_{(0)}\ \mathbf{L}_{(0)}$.

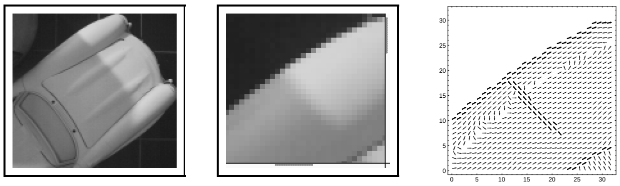

Figure 5: Left: this test image provide an example of shape recovery for a real image; middle: the image used is a close-up of the fender of a car; right: the corresponding edge map and shading flow field.

The linearity of Lambert's model allows the definition of an equivalent light source as the sum of the light source visible from the surface. Hence, the lighting condition on each side of the shadow boundary can be described by $\lambda\ \mathbf{L}$ and $\lambda_s\ \mathbf{L}_s$, the latter corresponding to the shadow side:

$$\lambda\ \mathbf{L} = \sum_i \lambda_{(i)}\ \mathbf{L}_{(i)} \ , \qquad \lambda_s\ \mathbf{L}_s = \sum_{i,\ i \neq 0} \lambda_{(i)}\ \mathbf{L}_{(i)} \ ,$$

hence this simple relationship:

$$\lambda_s\ \mathbf{L}_s = \lambda\ \mathbf{L} - \lambda_{(0)}\ \mathbf{L}_{(0)} \ .$$

At a point on the boundary, the onset of the occlusion can be caused either by a possibly distant opaque object or by the surface itself slanting away from the light source. We call these two cases, respectively, a cast shadow boundary and an attached shadow boundary. These account for different image structures.

To check if an image curve corresponding to an image irradiance (or intensity) discontinuity occurs at the shadow boundary, we look at the difference in image irradiance as we enter in the shadow: $E - E_s = \rho\ \lambda_{(0)}\ \mathbf{L}_{(0)} \cdot \mathbf{N}$. For a cast shadow, $\mathbf{L}_{(0)} \cdot \mathbf{N} > 0$, hence $E_s < E$. As our intuition suggests, there is a discontinuity. For an attached shadow, $\mathbf{L}_{(0)} \cdot \mathbf{N} = 0$; the incoming rays from the occluded source only graze the surface, hence $E_s = E$. Even though there is a marked decrease in surface illumination, the image irradiance is continuous.

To check if there is a shading flow field discontinuity at the shadow boundary, we compare the shading flow orientations on the both sides of a shadow boundary:

$$\tan\theta = \frac{\mathbf{L}\cdot\frac{\partial}{\partial y}\mathbf{N}}{\mathbf{L}\cdot\frac{\partial}{\partial x}\mathbf{N}}, \quad \tan\theta_s = \frac{\mathbf{L}_s\cdot\frac{\partial}{\partial y}\mathbf{N}}{\mathbf{L}_s\cdot\frac{\partial}{\partial x}\mathbf{N}}. \quad (4)$$

Since for cast and attached shadows, $\mathbf{L}$ and $\mathbf{L}_s$ are independent, the shading flow orientation on each side of the shadow boundary differs in general. One should note that in an image of a scene, coincidental alignments can occur at some points along the boundaries.

Observation 4.2 *For a generic scene, a cast shadow boundary is marked by a discontinuity in the image irradiance[3] and by a discontinuity in the shading flow field.*

Observation 4.3 *For a generic scene, an attached shadow boundary is marked by a continuous image irradiance and by a discontinuity in the shading flow field.*

Thus, the signature of the cast and the attached shadow boundaries differ by their image irradiance discontinuity. The signature of the attached shadow boundary can be seen across the sphere in Fig. 3.

Changes in reflectivity occur when the material properties change. If we consider a surface for which the reflectivity changes abruptly, the change will occur along a curve. Let ρ_l and ρ_d be respectively the albedo on the lighter and darker side of that curve.

Now, we examine the behaviour of the image irradiance on both sides of the curve where the reflectivity changes. If the image irradiance on the lighter side is denoted as E_l, then the image irradiance on the darker side is given by $E_d = \frac{\rho_d}{\rho_l} E_l$. If the surface is illuminated, then $E_d < E_l$ since $\rho_d < \rho_l$. Hence, another confirmation of an intuitive result.

The shading flow field has a different behaviour. Since it is independent of the albedo, $\tan\theta_d = \tan\theta_l = \frac{\mathbf{L}\cdot\frac{\partial}{\partial y}\mathbf{N}}{\mathbf{L}\cdot\frac{\partial}{\partial x}\mathbf{N}}$ the orientation of the shading flow is continuous even where an abrupt change in reflectivity occurs.

Observation 4.4 *An abrupt change in reflectivity is marked by a discontinuity in the image irradiance and a continuous shading flow field.*

[3]Strickly speaking, a discontinuity occurs only if the hidden light is a point source and if there is no atmospheric dispersion. The size of the image curve detectors defines a notion of scale. This scale relaxes the previously stated condition to allow light sources which are somewhat spatially extended and some limited amount of blurring due to atmospheric dispersion.

The signature of the abrupt change in reflectivity can be seen across the sphere and the plane in Fig. 1.

The difference in the image irradiance at surface discontinuities is given by the following $E - E' = \rho\,\lambda\,\mathbf{L}\cdot(\mathbf{N}-\mathbf{N}')$. The surface orientations on each side of a surface discontinuity ($\mathbf{N}$, $\mathbf{N}'$, respectively) are not related, hence $\mathbf{N}-\mathbf{N}'$ is usually not equal to zero. Since the illumination direction is not related to $\mathbf{N}-\mathbf{N}'$, the dot product $\mathbf{L}\cdot(\mathbf{N}-\mathbf{N}')$ is also usually not equal to zero.

Note that the occurrence of a point, somewhere along an occlusion, for which $\mathbf{L}\cdot(\mathbf{N}-\mathbf{N}') = 0$ is expected. For a generic situation, this point marks an inversion of contrast; Fig. 3 provides two examples of such points.

For the shading flow field, the consideration that the surface orientations (and its derivatives) on each side of a surface discontinuity are not related, yields the following observation.

Observation 4.5 *Surface discontinuities are marked (for a generic scene) by discontinuities in both the image irradiance and the shading flow field.*

The signature of surface discontinuities can be seen all around the spheres in Figs. 1 and 3, and where the fender meets the hood in Fig. 5.

Surface discontinuities occur when one surface partially occludes another. In such a case, when the albedos of the two surfaces are different, a change in albedo occurs at the same image location as the occlusion. The same example can be reworked for lighting conditions; in this case, the front surface could occlude some light sources from the back surface. Again, a change in lighting condition occurs at the same image location as the occlusion. These are not due to a coincidental alignment, these are rather due to a common cause.

By bringing the previous observations together, we get the beginning of an answer to our opening dilemma (Table 1).

Observation 4.6 *The different discontinuity signatures allow us to distinguish occluding boundaries from abrupt albedo changes or from attached shadow boundaries.*

For generic surfaces, there is only one ambiguity that cannot be resolved prior to shape estimation. Discontinuous surfaces and cast shadow boundaries on a smooth surface cannot be distinguished. The interpretation is nevertheless trivial once the surface shape is known (which is determined a shape from shading algorithm, see [3]).

disc. int.	disc. $\vec{S}$	characteristic	probable interpretation		
			surface	albedo	lighting
no	no	—	**smooth**	**constant**	**constant**
		parallel flow field	smooth parabolic	constant	shadow boundary (attached)
no	yes	—	**smooth**	**constant**	**shadow boundary (attached)**
yes	no	—	**smooth**	**abrupt change**	**constant**
		parallel flow field	smooth parabolic	constant	shadow boundary (cast)
		parallel flow field	aligned parabolic surfaces junction	—	—
yes	yes	—	**smooth**	**constant**	**shadow boundary (cast)**
		—	**discontinuous**	—	—

Table 1: This table shows the relations between image discontinuities and scene discontinuities. The shading flow field $\vec{S}(x,y)$ is perpendicular to the normalized gradient intensity $\frac{\vec{\nabla} I}{|\vec{\nabla} I|}$.

If we consider non-generic surfaces such as parabolic or planar surfaces[4], more ambiguities are introduced. However, these non-generic surfaces exhibit characteristic shading flow fields (parallel or null flow fields) that allow us to sort them out.

Parabolic surfaces exhibit parallel shading flow fields and, furthermore, the orientation of the field is only dependent on the geometric aspect of the surface: this leads to the following consequences.

- Although the attached shadow boundary no longer exhibits its shading flow field discontinuity, the magnitude of the gradient remains discontinuous along a line parallel to the flow. It is thus possible to distinguish it from a smooth surface of uniform albedo under constant illumination condition.

- The cast shadow boundary and discontinuous aligned parabolic surfaces no longer exhibit a discontinuity in the shading flow field. But again the magnitude of the gradient remains discontinuous along a line, this time parallel to an intensity discontinuity which, for the cast shadow boundary's case, is not necessarily parallel to the shading flow field.

[4]A surface is said to be *parabolic* if one of its principal curvatures is null; a surface is said to be *planar* if both principal curvatures are null.

Planar surfaces always exhibit null shading flow fields, hence they also generate ambiguities. It is not possible to distinguish between an abrupt change in albedo, a shadow boundary[5], and a surface discontinuity where only the surface orientation changes; but see Langer [10].

The ambiguities introduced by parabolic surfaces do not completely prevent the interpretation of the geometric structures of the image. There is always some characteristic shading flow field where these ambiguities occur. Our analysis can be carried everywhere else.

At least one such ambiguity can be resolved. An attached shadow boundary on a smooth parabolic surface of constant albedo does not demarcate itself from a smooth surface of constant albedo under constant lighting condition, given the image geometrical structures identified thus far. We examine the behaviour of the image irradiance gradient magnitude near an attached shadow boundary. Since $\lambda \mathbf{L} \geq \lambda_{(0)} \mathbf{L}_{(0)}$ and $\lambda_{(0)} \mathbf{L}_{(0)} > 0$, the magnitude of the image irradiance differs from each side of the shadow boundary.

Observation 4.7 *An attached shadow boundary is marked by a discontinuity in the image irradiance gradient magnitude (even for a parabolic surface).*

[5]Only the cast shadow boundary can occur on a flat surface, an attached shadow boundary involves a surface curving away from a light source.

disc. $\left\|\vec{\nabla} I\right\|$	disc. int.	disc. $\vec{S}$	characteristic	probable interpretation surface	albedo	lighting
no	no	no	—	**smooth**	**constant**	**constant**
yes	no	no	parallel flow field	smooth parabolic	constant	shadow boundary (attached)

Table 2: This table shows other relations between image discontinuities and scene discontinuities.

The discontinuity of the gradient magnitude lifts the ambiguity as shown in Table 2. Image curves capturing this type of discontinuity can also be extracted from the image using logical/linear operators [7]).

We have shown that much can be inferred about the lighting conditions and the object in the scene a priori to the computation of shape from shading. We conclude this section by stressing the usefulness of the geometrical structures of the image — the shading flow field and the tangent field used together, not only on their own — for the understanding of the scene.

Acknowledgments

This work is supported by AFOSR, NSERC, IRIS Centre of Excellence.

References

[1] Martin Bichsel and Alex P. Pentland. A simple algorithm for shape from shading. In *Proceedings, Conference on Computer Vision and Pattern Recognition*, pages 459–465, Champaign, Illinois, June 15–18 1992. IEEE-CS, IEEE Computer Society Press.

[2] Andrew Blake. Improbable views. *Nature*, 368:498–499, April 1994.

[3] Pierre Breton. *Recovering Shape from Shading and Contour.* PhD thesis, McGill University, Montréal, November 1994.

[4] William T. Freeman. The generic viewpoint assumption in a framework for visual perception. *Nature*, 368:542–545, April 1994.

[5] Berthold K.P. Horn and Michael Brooks. *Shape from Shading.* MIT Press, Cambridge, Massachusetts, 1989.

[6] R. A. Hummel and S. W. Zucker. On the foundation of relaxation labeling processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 5:267–287, 1983.

[7] Lee A. Iverson and Steven W. Zucker. Logical/linear operators for image curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1995.

[8] J. J. Kœnderink and A. J. van Doorn. Photometric invariants related to solid shape. *Optica acta*, 27(7):981–996, 1980.

[9] J. J. Kœnderink and A. J. van Doorn. Representation of local geometry in the visual system. *Biological Cybernetics*, 63:291–297, 1987.

[10] M. S. Langer and S. W. Zucker. Shape from shading on a cloudy day. *Journal of the Optical Society of America*, February 1994.

[11] J. Oliensis and P. Dupuis. A global algorithm for shape from shading. In *Proceedings of the 4th International Conference on Computer Vision*, pages 692–701, Berlin, Germany, 1993.