Contours, Shading flows, and Generalized Shape from Shading

Pierre Breton Steven W. Zucker

Department of Electrical Engineering
McGill University
Montréal (Québec) H3A 2A7, Canada
e-mail: breton@cim.mcgill.ca

Abstract

The classical approach to shape from shading problems is to find a numerical solution of the image irradiance partial differential equation (PDE). It is usually assumed that the parameters of this equation (the light source direction and surface albedo) can be estimated in advance. For images which contain shadows and occluding contours, this decoupling of problems is artificial. In this paper, we redefine the shape from shading problem and we show that it is necessary to use the image geometric structures to address issues related to discontinuities.

1 Introduction

Shape from shading is a classical problem. Ernst Mach (1866) was perhaps the first to establish a formal relationship between image [10] and scene domains, and to capture their inter-relationships in a PDE. Horn set the modern approach by focusing on the solution of these PDEs by classical and numerical techniques [4–6], and to do so, he introduced the *image illumination equation*:

$$E(x,y) \;=\; R(p,q) \;,$$

where p and q are partial derivatives of the surface with respect to x and y. The basic assumption here is that all variations in the image irradiance are due to variations in the surface gradient, thus to variations in surface orientation.

Problem 1. CLASSICAL SHAPE FROM SHADING *Assuming that the scene is all the same reflectivity; the scene is illuminated by a single distant light source; the scene is composed of a single smooth surface; the surface is matte; the image is formed by orthographic projection. Given the albedo ρ; the illumination λ; the direction of the light source $\mathbf{L}$; the image irradiance $E(x,y)$. Recover the surface shape.*

The matte surface is traditionally modeled with Lambert's reflectance function so the image irradiance[1] equation becomes

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

where $I(x,y)$ is the image intensity 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).

Fundamental to the classical shape from shading problem is:

Assumption 1.1 (Horn) *Variations in image intensity are entirely due to variations in surface orientation.*

Previous research on shape from shading concentrated on images such as Fig. 1, i.e., images of a smooth surface of constant albedo illuminated by a single distant light source under Lambert's reflectance model. For such images, it is tempting to decouple the surface estimation problem from the light source estimation problem and the surface albedo estimation problem. However, this is a major drawback, because such a decoupling of problems yields algorithms whose domain of application is very limited. To illustrate, such algorithms could not deal with: images of scenes with shadows, i.e. parts are illuminated differently; images of scenes in which surfaces have different albedos (or reflectance coefficients), i.e. the scene is not entirely the same colour; images of scenes with surface geometric discontinuities, i.e. in which the scene contains multiple surfaces either abutting or not. In fact, it is quite common to encounter such images, and it is quite rare to encounter the classical ones.

[1] The image irradiance $E(x,y)$ is often replaced by the image intensity $I(x,y)$. One assumes a linear relation between the image irradiance $E(x,y)$ and the image intensity $I(x,y)$.

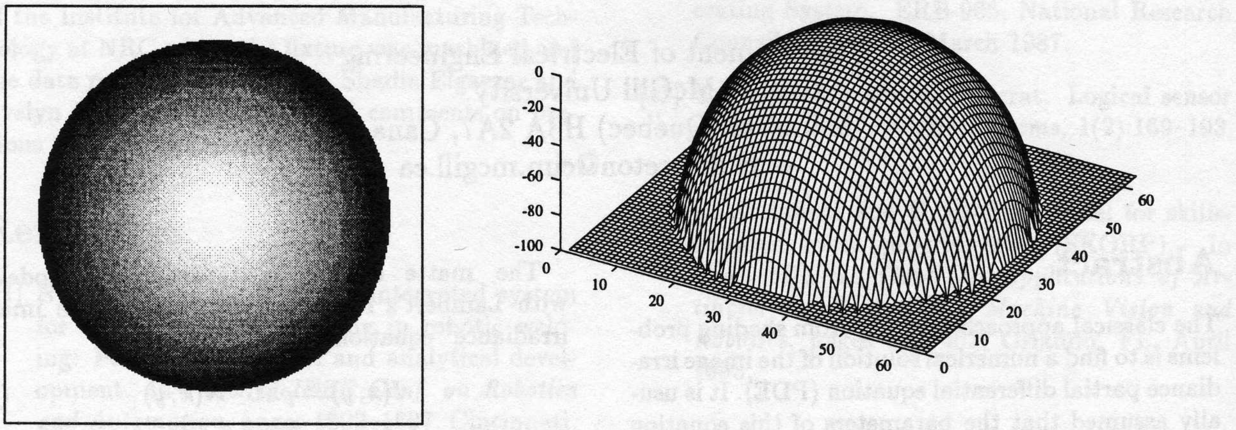

(a) (b)

Figure 1: We successfully recover a constant albedo matte scene illuminated by a single distant point light source using the classical shape from shading assumptions. An ideal intensity image of (a) and the corresponding recovered depth map (b) are shown here.

1.1 Analysis of a "Classical" Shape from Shading Algorithm

Recently, new solutions of the "classical" shape from shading problem were published [1,9,12]. We'll use Bichsel and Pentland's (B&P) algorithm [1] to show examples of classical shape from shading[2] (see Fig. 1).

In Fig. 2(a), we show an image of the sphere in front of a plane but part of plane and sphere is in a darker colour than the rest. The depth map obtained with B&P's algorithm is shown in Fig. 2(b). We can observe a deformation of the sphere. Similar deformations was observed for shadows (Fig. 3). Although the sphere seems like an extremely simple example, this discussion illustrates how it is sufficiently complex to demonstrate many subtleties in shape from shading.

[2]The two key subroutines were copied directly from their conference paper. However, it was necessary to write code for the input/output as well as the calls for these subroutines.

2 The General Shape from Shading Problem

The assumptions underlying the classical shape from shading problem are too restrictive. A human observer confronted with a static, monocular view of a scene will succeed in obtaining some estimate of the shapes of the surfaces within it, even when some of the classical setting's constraints (single smooth surface, single light source, no shadow, single albedo scene) are relaxed (multiple smooth surfaces, multiple light sources, shadows, multiple albedo scene). Thus the classical constraints can be relaxed in principle; but how far?

We redefine the shape from shading problem to retain only the essence of the basic assumption 1.1, i.e.:

Assumption 2.1 *For the generalized shape from shading problem,* smooth *variations in intensity are entirely due to* smooth *variations in surface orientation.*

To make the distinction clear, we emphasize the fact that this assumption is only concerned with variations that are *smooth.* Abrupt variations of intensity can be due to abrupt variations in surface

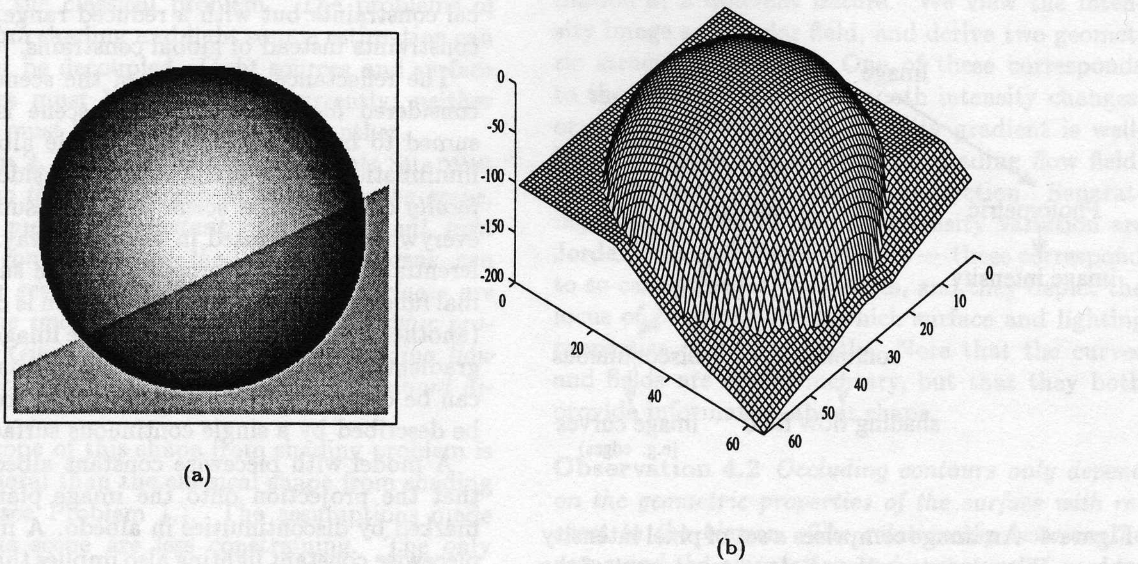

Figure 2: Using B&P's algorithm, we attempt to recover a simple scene in which the surface albedo changes abruptly — the scene is not entirely the same colour. The intensity image (a) and the corresponding recovered depth map (b) are shown in this figure. This result shows clearly that the algorithm based on the classical shape from shading assumptions fails to accurately recover a multiple albedo scene.

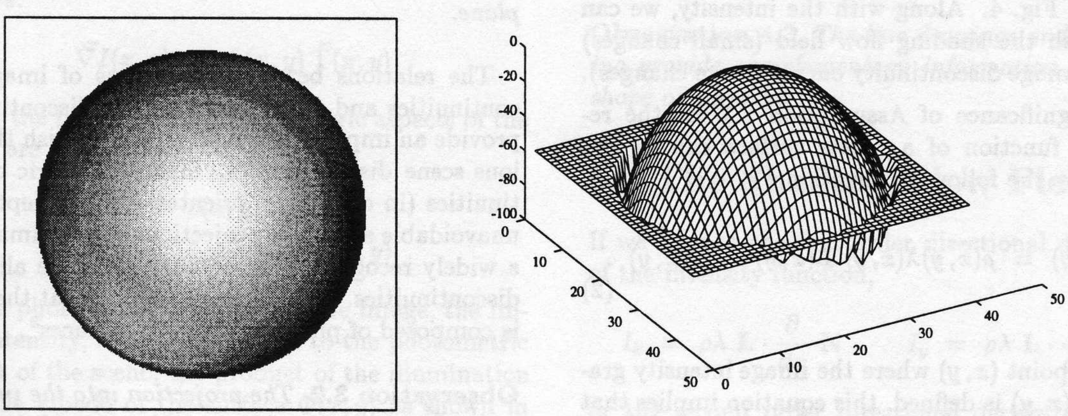

Figure 3: Another example of shape recovery attempt with B&P algorithm. (a) Parts of the sphere are illuminated with only one light source — the upper right and the lower left parts — and the central part of the sphere and the back plane are illuminated by both light sources. (b) the recovered depth map erroneously shows a "valley" to account for the shading variability around the edge.

orientation, but they can also be due to abrupt variations in the surface albedo or abrupt variations in the lighting conditions. We examine such abrupt changes later.

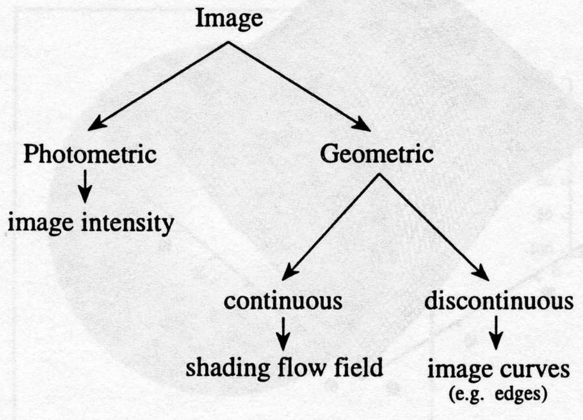

Figure 4: An image comprises a set of pixel intensity values. These capture the photometric aspect of the image. An image also implicitly defines geometric structures. These emerge from the neighbourhood relationships of the image pixels. These structures can describe either continuous properties of the image or discontinuous ones.

The changes of intensity occurring at various scales reveal that there is more to an image than just the photometric aspect. The various structures are shown in Fig. 4. Along with the intensity, we can distinguish the shading flow field (small changes) and the image discontinuity curves (large changes).

The significance of Assumption 2.1 for the reflectance function of a matte surface can be expressed by the following equation:

$$\nabla I(x, y) = \rho(x, y)\lambda(x, y)\ \mathbf{L}(x, y) \cdot \nabla\mathbf{N}(x, y)\ . \tag{2}$$

At every point (x, y) where the image intensity gradient $\nabla I(x, y)$ is defined, this equation implies that the surface albedo $\rho(x, y)$, the illumination $\lambda(x, y)$ and the direction of the illuminant $\mathbf{L}(x, y)$ are *locally* constant. It also implies the existence of a differentiable surface normal $\nabla\mathbf{N}(x, y)$. Note that the other terms in the total derivative ∇I are assumed nonexistent according to the above comments about scale; i.e., $\nabla\mathbf{L} = 0$, $\nabla\rho = 0$, and $\nabla\lambda = 0$. We elaborate on the conceptual differences in the following sections.

3 Local vs. Global Scene Constraints

The key to our approach is to reconsider the classical constraints but with a reduced range, i.e. local constraints instead of global constrains.

The reflectance properties of the scene are only considered *locally* constant; the scene is not presumed to be described with a single albedo. The illumination of the scene is also considered to be *locally* constant; the scene is not presumed to be everywhere illuminated in the same way. The differentiability (a local property) of the surface normal function implies that this function is continuous (another local property). Where the image intensity gradient is not defined, the surface normal function can be discontinuous; the scene is not presumed to be described by a single continuous surface.

A model with piecewise constant albedo implies that the projection onto the image plane will be marked by discontinuities in albedo. A model with piecewise constant lighting also implies that the projection onto the image plane will be marked by discontinuities in illumination and illuminant direction. Similarly, a model with piecewise smooth surfaces implies that the projection onto the image plane will be marked by discontinuities in surface depth, orientation and shape. All of these discontinuities form contours.

Observation 3.1 *The discontinuities in the scene are marked by curves of discontinuity in the image plane.*

The relations between the curves of image discontinuities and the curves of scene discontinuities provide an important insight to distinguish the various scene discontinuities. Since geometric discontinuities (in curvature, orientation, and depth) are unavoidable and their projection into the image has a widely recognized importance [2, 7], we allow for discontinuities. We therefore assume that the scene is composed of piecewise smooth surfaces[3].

Observation 3.2 *The projection into the image of geometric discontinuities provides an additional cue to help shape recovery.*

Since we are considering local constraints for the generalized shape from shading problem, the vector $\rho\lambda\mathbf{L}$ is constant only within some neighbourhood.

[3]Since the reflectance function depends on the existence of a differentiable surface normal, allowing surfaces that are nowhere smooth is clearly inappropriate.

One should typically expect this vector to take multiple values for any given scene. The problem is more complex and one cannot expect to extract the vector $\rho\lambda\mathbf{L}$ based on the statistical considerations used for the classical problem. The problems of shape from shading and light source estimation can no longer be decoupled. Light sources and surface properties must be handled concurrently; neither problem must be solved "before" the other.

Problem 2. Generalized Shape from Shading *Assuming that colours can be present in the scene, but are piecewise constant; several distant light sources can illuminate the scene; the scene can comprise several smooth surfaces; the surfaces are matte; the image is formed by an orthographic projection. Given the tangent field; the shading flow field. Recover the surface shape; the illuminant direction.*

The scope of this shape from shading problem is more general than the classical shape from shading setting (see Problem 1.). The assumptions made about the scene are less constraining. The only given variables are the geometric structures of the image — these can be extracted from the image intensities.

4 Geometric vs. Photometric Image Structures

To emphasize the dichotomy between the photometric and geometric structures, we rewrite Eq. 2 as follows:

$$\vec{\nabla} I(x,y) \;=\; \Phi(x,y)\;\vec{\Gamma}(x,y)$$

where the photometric and geometric aspects of the scene are respectively:

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

The photometric structure of the image, the image intensity, is directly related to the photometric aspect of the scene, the product of the illumination and the albedo of the surface $\Phi(x,y)$ as shown in Eq. 1. Since we are letting both the illumination and the albedo be functions of position, the image intensity is a poor choice for initial data of the shape from shading problem. The geometric information about the scene would be confounded with the photometric.

Observation 4.1 *The relaxation of the scene constraints leads to the rejection of the photometric structure of the image as initial data for the shape from shading problem.*

The geometric aspect of the image reveals information of a different nature. We view the intensity image as a scalar field, and derive two geometric structures from it. One of these corresponds to the regions depicting smooth intensity changes, or regions in which the intensity gradient is well-behaved. This will become the shading flow field, and it is developed later in this section. Separating these regions of smooth intensity variation are Jordan curves of discontinuities — these correspond to so-called "edges" in images, and they depict the locus of positions along which surface and lighting properties change abruptly. Note that the curves and fields are complementary, but that they both provide information about shape.

Observation 4.2 *Occluding contours only depend on the geometric properties of the surface with respect to the viewer. The relationship between the scene and the orientation and curvature of an edge element is therefore independent of the the illumination model and the surface reflectance model.*

At those positions where the image intensity gradient is not defined, image curves can provide precious information. But if we provide the shading information along with the line drawing, we can notice details that we have missed previously. The shading provides evidence of undulations that the line-drawing cannot capture.

Observation 4.3 *The line drawings and the shading provide complementary information about the shape of surfaces.*

5 The Shading Flow Field

If we look at the first order directional derivatives of the intensity function,

$$I_x = \rho\lambda\,\mathbf{L}\cdot\frac{\partial}{\partial x}\mathbf{N}\ ,\quad I_y = \rho\lambda\,\mathbf{L}\cdot\frac{\partial}{\partial y}\mathbf{N}\ ,$$

or the second order directional derivatives of the intensity function,

$$\begin{aligned}I_{xx} &= \rho\lambda\,\mathbf{L}\cdot\frac{\partial^2}{\partial x^2}\mathbf{N}\ ,\quad I_{yy} = \rho\lambda\,\mathbf{L}\cdot\frac{\partial^2}{\partial y^2}\mathbf{N}\ ,\\ I_{xy} &= \rho\lambda\,\mathbf{L}\cdot\frac{\partial^2}{\partial x\partial y}\mathbf{N}\ ,\end{aligned}$$

we note that all these partial derivatives depend on the variable illumination and albedo. Again, as in

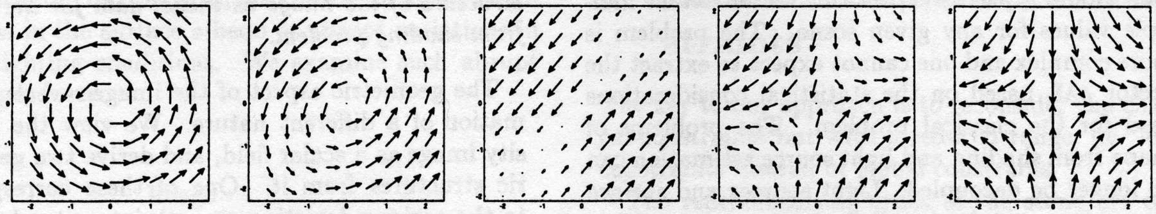

Figure 5: This sequence shows flow fields for which the distortion vector differs. The left most depicts a pure curl, the right most depicts a pure deformation, and in between, the two components blend with different proportions. Whereas the curl is symmetric with respect to a point, the deformation is symmetric with respect to an axis. Note how the nature of the singularities differs as the distortion changes.

the image intensities, these measures confound the scene's geometric information with the scene's photometric properties. Therefore, they are not appropriate to solve the generalized shape from shading problem.

The directional derivatives' information can nevertheless be combined in ways that are independent from the variable illumination and albedo. One such way is the *shading flow field.*

Definition 5.1 *The* shading flow field *is the vector field which indicates the direction in which the image intensity remains constant. It is therefore perpendicular to the image intensity gradient field.*

Although the following calculations are elementary, their implications are deep.

Theorem 5.1 *For the generalized 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 and the lighting with respect to the viewer.*

Proof: The orientation of the gradient field θ is related to a ratio of directional derivatives: $\tan\theta = I_y/I_x$. Since both directional derivatives are directly proportional to the product of the illumination and the albedo,

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

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

Corollary 5.1 *The shading flow field (where it is defined) is independent from 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 5.2 *The curvatures of both the shading flow field and the normalized gradient field are also independent from 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:

$$\kappa_s = \frac{2I_xI_yI_{xy} - I_x^2I_{yy} - I_y^2I_{xx}}{(I_x^2+I_y^2)^{\frac{3}{2}}},$$

$$\kappa_g = \frac{I_xI_y\left(I_{xx}-I_{yy}\right) - I_{xy}\left(I_x^2-I_y^2\right)}{(I_x^2+I_y^2)^{\frac{3}{2}}}$$

Since the first and second directional derivatives are directly proportional to the product of the illumination and the albedo, in both instances, the numerator and the dominator would be proportional the the cube of this product and cancel out. We define a third "curvature-like" quantity, $\kappa_c = \left(2I_xI_yI_{xy} - I_x^2I_{yy} - I_y^2I_{xx}\right)\left(I_x^2+I_y^2\right)^{-\frac{3}{2}}$ to express the variation of the shading flow field in terms of a differential invariants (see Fig. 5):

$$\begin{aligned} \operatorname{def}\vec{S}\,\sin(2\varphi) &= -\kappa_s + \kappa_c\ , \\ \operatorname{def}\vec{S}\,\cos(2\varphi) &= -2\kappa_g\ , \\ \operatorname{curl}\vec{S} &= -\kappa_s - \kappa_c\ . \end{aligned}$$

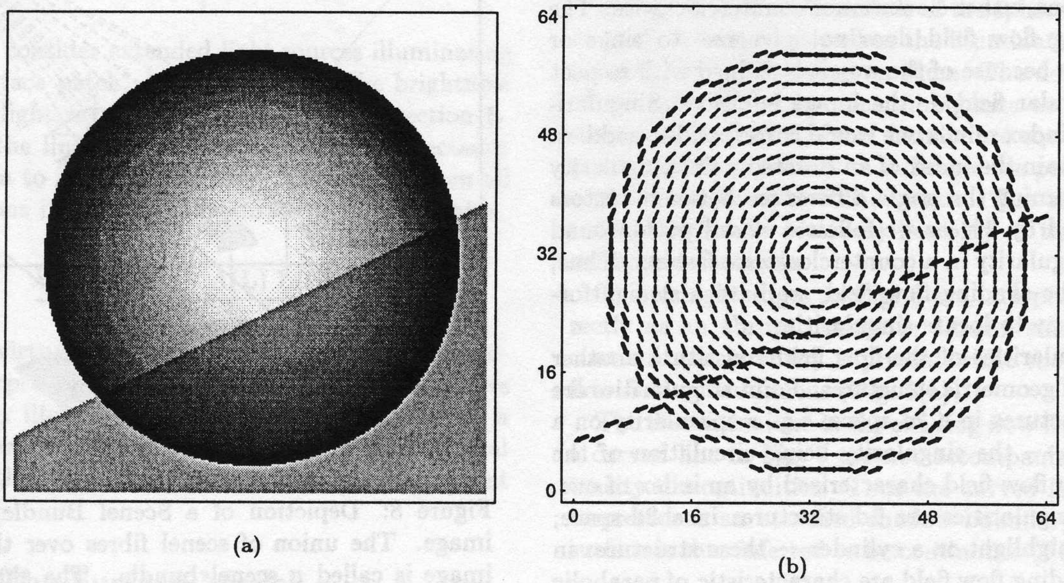

Figure 6: (a) An ideal image of a simple scene illuminated by a single distant point light source but with an abrupt change in albedo and (b) the corresponding shading flow field.

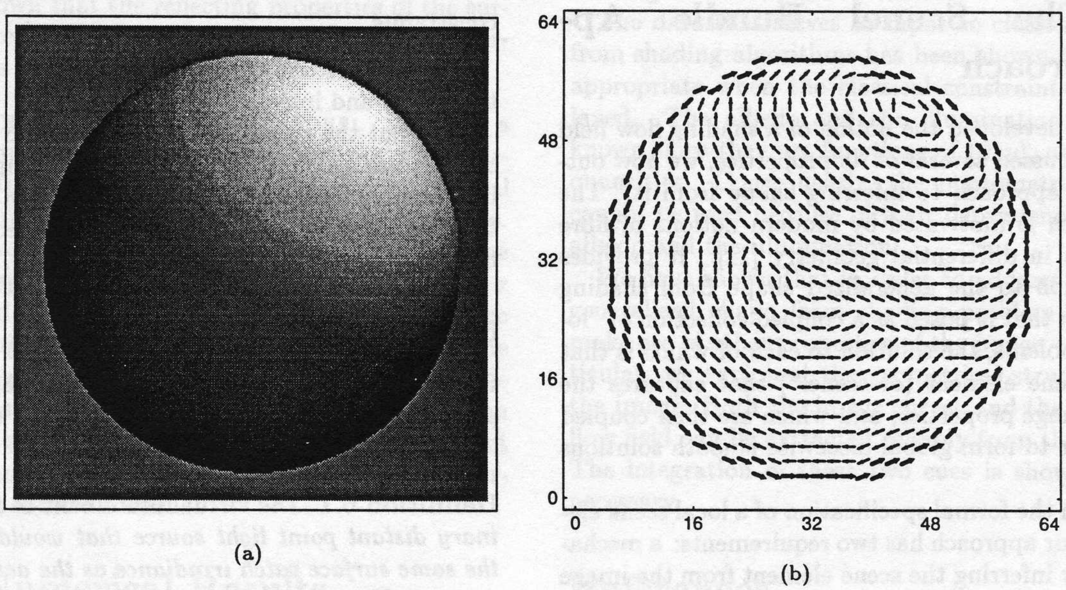

Figure 7: (a) An ideal image of a simple scene illuminated by two distant point light sources and (b) the corresponding shading flow field.

Figures 6 and 7 provide two examples of shading flow fields.

The shading flow field can exhibit singularities. Singularities of index one would correspond to circulations, either clockwise or counterclockwise. The shading flow field does not give rise to sinks or sources because of the way it is defined with respect to a scalar field — the image intensity. Singularities of index minus one would correspond to saddles. We obtain the index of an isolated point singularity by summing the angle differences between vectors (divided by 2π) as we follow a closed path around the singularity in a counterclockwise fashion. Thus, from the shading flow field, we have a straightforward way to locate singularities.

Singularities of the flow field constitute another class of geometric structures. Point singularities are 0d-structures in a 2d-space; e.g. a specularity on a surface — the singularity being circulation of the shading flow field characterized by an index of one. Line singularities are 1d-structures in a 2d-space; e.g. a highlight on a cylinder — these structures in the shading flow field are characteristic of parabolic surface patch (or line [8]). Undefined regions are 2d-structures in a 2d-space; e.g. the luminance of a flat surface — characteristic of planar surface patch, the intensity is constant thus the gradient is null.

6 The Scenel Bundle Approach

Having developed the notion of a shading flow field and discussed several of its properties, we now outline an approach to inferring shape from it. The approach is motivated by modern notions of fibre bundles in differential geometry [11]. It provides a solution to the generalized shape from shading problem that is posed as a coupled collection of "local" problems, the solution to each of which is that local scene element (or scenel[4]) that captures the local image properties, and which are then coupled together to form global piecewise smooth solutions (Fig. 8).

Given the formal specification of a local scene element, our approach has two requirements: a mechanism for inferring the scene element from the image (or more precisely, from the image geometric structures); a mechanism for the local to global transition, i.e. for sewing the local scene patches together in a consistent fashion. We have already described the outline of our approach in a previous conference [3]. We will only mention insightful details of the model and interpret our results.

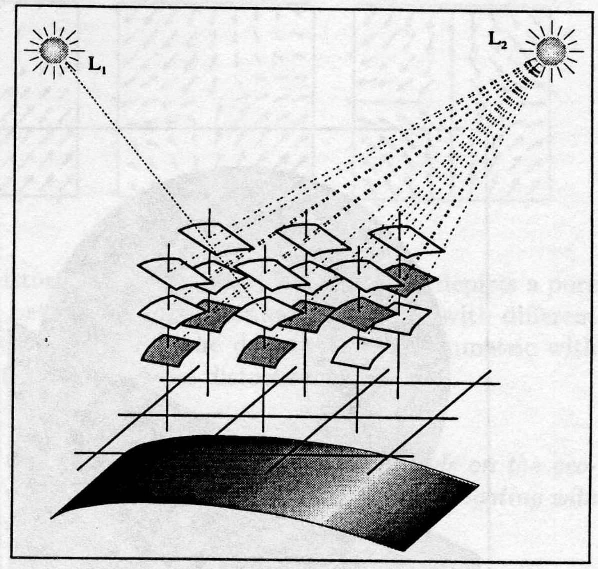

Figure 8: Depiction of a Scenel Bundle over an image. The union of scenel fibres over the entire image is called a scenel bundle. The shape from shading problem is formulated as determining sections through the scenel bundle. Such a section is depicted by the shaded scenels, and represents a horizontal slice across the bundle. Scenel participation in a horizontal section is governed by surface smoothness and material and light source constancy constraints.

We want the *scene element* or *scenel* to be a local representation of the scene which comprises an illumination model and a surface model. In this section, we describe the information captured by the illumination model and the surface model.

6.1 The Illumination Model

Restricting our shading analysis to the matte component of the reflection allows us to use a simple local illumination model.

Definition 6.1 *The* virtual illuminant *is an imaginary distant point light source that would provide the same surface patch irradiance as the actual light sources. It comprises two attributes, λ and $\mathbf{L}$, that can be derived considering the linear property of Lambert's reflectance function.*

Consider a set of M distant point sources that illuminate the surface patch. Let this set be described by $\{\lambda_{(i)}\mathbf{L}_{(i)} : 1 \leq i \leq M\}$ where $\lambda_{(i)}$ and $\mathbf{L}_{(i)}$ are respectively the intensity and the direction

[4] *cf.* Pixel, voxel, ... scenel.

of the individual point light source. The attributes of the virtual illuminant are given by:

$$\lambda \, \mathbf{L} \equiv \sum_i \lambda_{(i)} \, \mathbf{L}_{(i)} ,$$

Now consider extended light sources illuminating the surface patch. Let $\Lambda(\mathbf{L})$ denote the brightness of the light ray that is incident from direction $\mathbf{L}$. Using the linear property of Lambert's reflectance function to integrate the light rays coming from all directions in the visible hemisphere $\mathcal{H}$, we obtain:

$$\lambda \, \mathbf{L} \equiv \frac{1}{\pi} \int_{\mathcal{H}} \Lambda(\mathbf{L}) \, \mathbf{L}(\Omega) \, d\Omega \, .$$

The virtual illuminant thus provides a very convenient representation of the illumination as it allows complex illumination to be described simply by a scalar and a unit vector. It is a much more general model than that of a simple infinitely distant point source.

6.2 The Surface Model

The surface model provides a local description of what reflects the light. We consider two distinct attributes of the surface model:

1. MATERIAL PROPERTIES: We have already shown that the reflecting properties of the surface are not related to the geometrical properties of the image, thus we ignore them.

2. SURFACE SHAPE DESCRIPTORS: We use a shape representation that captures the information of the first and second fundamental forms' coefficients to describe surface patches. Two principal curvatures (κ_1, κ_2) describe the shape up to rotation. Two angles, slant σ and tilt τ, are needed to describe the surface tangent plane orientation with respect to the viewer's coordinate frame. An additional angle ϕ is needed to describe the principal direction of the Darboux frame in the surface tangent plane. Cartan equations constrain neighbouring surface model.

7 Numerical Results

In Figs. 6 and 7, the curved edges that form a circle carry two interpretations since discontinuities in both the intensity and the shading flow field are present: either they mark a cast shadow boundary; or they mark a surface discontinuity. While the artificial scene is not meant to be characterized by a change in lighting condition, such interpretation is not incompatible with the image. The second interpretation is the one which we intend to exploit as these edges mark in fact an occluding contour.

In Fig. 9, the visible surface of the spheres are accurately recovered as well as the illumination directions of the light shining on them. The description of the back plane remains ambiguous as multiple cross-sections of the scenel bundle are equally supported. Each cross-section indicates the presence of a planar surface which is accurate. The ambiguity (which cannot be resolved) lies only in the planar surface's orientation.

The straight edges in Fig. 6 are interpreted correctly as an abrupt change in albedo, since a discontinuity in intensity is present while the shading flow field is continuous (Fig. 9(a)).

The discontinuities in the shading flow field along a curved line in Fig. 7 are not accompanied by intensity discontinuities. Thus the curved line is interpreted as an attached shadow boundary — the surface is thus inferred to be continuous (Fig. 9(b)).

8 Discussion

By relaxing the constraints of the classical shape from shading problem, we have given shape from shading an entirely new structure.

The data that serves as input to classical shape from shading algorithms has been shown to be inappropriate when the classical constraints are relaxed. The albedo and the illumination are not known since they are functions of (x, y), not global quantities. Consequently, the photometric values cannot be used because of their dependence on the albedo and the illumination.

We have shown that in order to address the new generalized shape from shading problem, we can make use of other properties of the image — in particular, we exploited the geometrical structures of the image. Both the image curves and the shading flow field can be extracted reliably from the image. The integration of these two cues is shown to be necessary.

References

[1] M. Bichsel and A. 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.

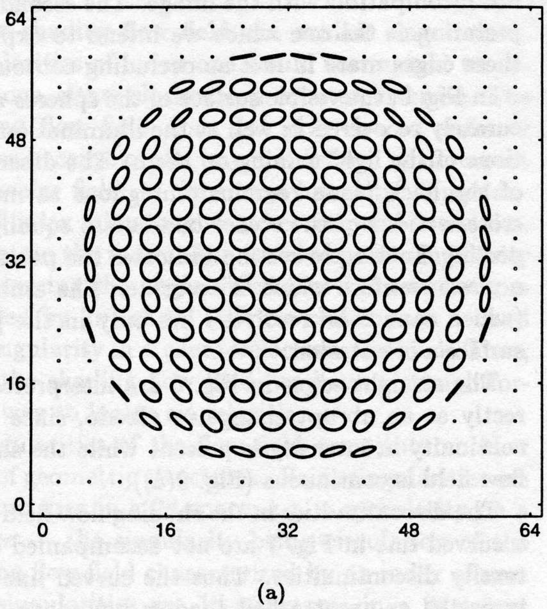

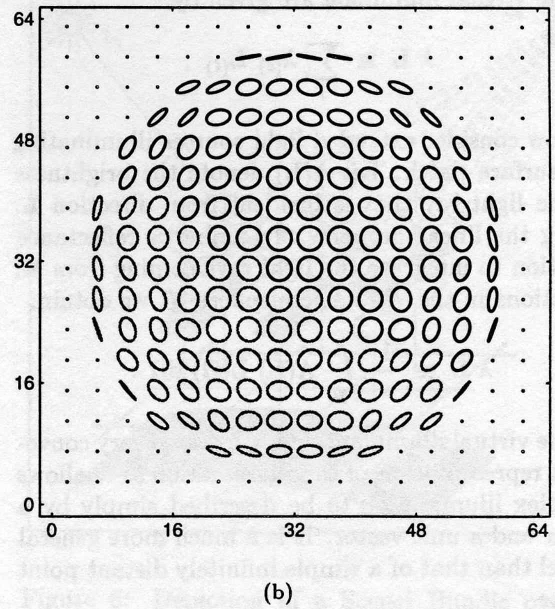

Figure 9: Using the generalized shape from shading assumptions, we recover (a) the scene shown in Figs. 6 and (b) the scene shown in 7 The difference between these results is not immediately visible since the surface orientations are exactly the same — it is only the light source orientation which differs. The shape recovery is illustrated by the surface normals. Different illuminant directions are coded with different gray levels.

[2] I. Biederman. Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94:115–147, 1987.

[3] P. Breton, L. A. Iverson, M. S. Langer, and S. W. Zucker. Shading flows and scenel bundles: A new approach to shape from shading. In *Computer Vision – ECCV 92*, pages 135–150, Santa Margherita Ligure, Italy, 19–22 May 1992. Springer-Verlag.

[4] B. K. P. Horn. *Shape from Shading: a method for obtaining the shape of a smooth opaque object from one view.* PhD thesis, Massachusetts Institute of Technology, June 1970.

[5] B. K. P. Horn and M. J. Brooks. The variational approach to shape from shading. *Computer Vision, Graphics, and Image Processing*, 33:174–208, 1986.

[6] K. Ikeuchi and B. K. P. Horn. Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, 17:141–184, 1981.

[7] J. J. Kœnderink. What does the occluding contour tell us about solid shape? *Perception*, 13:321–330, 1984.

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

[9] 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.

[10] F. Ratliff. *Mach Bands: Quantitative Studies on Neural Networks in the Retina.* Holden-Day, San Francisco, CA, 1965. p. 285.

[11] M. Spivak. *A Comprehensive Introduction to Differential Geometry.* Publish or Perish, Berkeley, 1979.

[12] R. Zhang, P.-S. Tsai, J. E. Cryer, and M. Shah. Analysis of shape from shading techniques. In *Proceedings, Conference on Computer Vision and Pattern Recognition*, pages 377–384, Seattle, Washington, June 21–23 1994. IEEE-CS, IEEE Computer Society Press.