VARIATIONAL TECHNIQUES AND OPEN PROBLEMS IN THE RECOVERY OF SURFACE SHAPE FROM IMAGE SHADING

M.J. Brooks

Abstract. Shading, or brightness variation, exhibited within an image is a well-known cue for what is variously termed surface shape, depth, or form. As examples, artists commonly use tonal values in drawings to convey a realistic impression of 3-dimensionality, and make-up is sometimes used to give facial shape a flattering appearance. Just how shape might be recovered from shading has long been studied. However, it is only with the advent of computer vision in the last 25 years that substantial progress has been made. It transpires that the problem takes the form of a first-order partial differential equation. The pioneering automated method for determining surface shape from image shading was based on the use of the equations of characteristic strips. This suffered from several drawbacks which were partly overcome with the introduction of parallel schemes, based on minimisation techniques, for use on a 2-dimensional rectangular grid. Typically, these schemes are directed at recovering components of surface normals defined over a region of the image. Recovery of actual depth is left to a subsequent stage of integration. The variational calculus has now become the standard mathematical device by which new schemes are developed. It has also served as a useful means by which their shortcomings may be analysed and rectified. This paper presents the major iterative techniques in the field, and discusses various open problems and difficulties that remain.

1. Introduction

A monochrome photograph of a smooth object will typically exhibit brightness variation, or *shading*. Of interest to researchers in computer vision has been the inverse problem of how object shape may be extracted from image shading. This *shape-from-shading* problem has been shown by Horn [7] to correspond to that of solving a first-order partial differential equation. Specifically, one seeks a function $u(x, y)$, representing surface depth in the direction of the z-axis, satisfying the *image irradiance equation*

$$R(u_x, u_y) = E(x, y)$$

over Ω. Here R is a known function (the so-called *reflectance map)* capturing the illumination and surface reflecting conditions, E is an image formed by (orthographic) projection of light along the z-axis onto a plane parallel to the xy-plane, and Ω is the image domain. Subscript denotes partial differentiation. In this formulation, it is implicitly assumed that:

- A small surface portion reflects light independently of its position in space. Thus, scene radiance emitted in a given direction is dependent only on the illumination, the light-scattering properties of the surface material, and the surface normal. By implication, light sources are infinitely far away, and internal surface reflections are disallowed.

- Image irradiance is equal to the projected scene radiance.

Subsequent to the elucidation of the problem's fundamental form, much effort has been expended on generating computational techniques for solving the image irradiance equation. It is the aim of this introductory article to summarise the main iterative techniques developed to solve this problem, and to indicate some remaining major difficulties in the field. For a more detailed treatment of many of the issues covered herein, the reader is referred to Horn and Brooks [9], as well as to the annotated bibliography in Horn and Brooks [10].

It is useful to consider a specific example of a shading problem. Assume that an overhead, distant point-source illuminates a *Lambertian* surface. A small portion of such a surface acts as a perfect diffuser, appearing equally bright from all directions. Let a small surface portion with normal direction $(-u_x, -u_y, 1)$ be illuminated by a distant, overhead point-source of unit power in direction $(0, 0, 1)$. The emitted radiance, as prescribed by Lambert's law, is given by the cosine of the angle between the two directions, namely $(u_x^2 + u_y^2 + 1)^{-1/2}$. Since image irradiance is assumed to be identical in value with scene radiance, the corresponding image $E(x, y)$ satisfies

$$(u_x^2 + u_y^2 + 1)^{-1/2} = E(x, y);$$

this, then, is the image irradiance equation appropriate for a Lambertian surface illuminated by a distant, overhead point-source, which may be written in the form of an *eikonal equation*

$$u_x^2 + u_y^2 = \mathcal{E}(x, y),$$

where $\mathcal{E}(x, y) = \big(E(x, y)\big)^{-2} - 1$. This relatively simple form is amenable to fundamental analysis (see Brooks et al. [1,2]). Now, in the event that our image is actually that of a unit sphere, our problem becomes that of solving

$$(u_x^2 + u_y^2 + 1)^{-1/2} = (1 - x^2 - y^2)^{1/2},$$

over the unit disc. If the light source is instead in the direction pointed to by the unit vector (a, b, c), then the image irradiance equation becomes

$$\frac{(-au_x - bu_y + c)}{(u_x^2 + u_y^2 + 1)^{1/2}} = ax + by + c(1 - x^2 - y^2)^{1/2},$$

over that portion of the image which is positively valued. It should be emphasised that these are the among the simplest cases arising in the shape-from-shading problem. Nevertheless, they arise from a reasonably realistic model that gives rise to very pleasing shading patterns. Often, however, the reflectance map may be more complex, or may be known only in numerical form.

It is important to appreciate the nature of the boundary conditions involved in solving the problem. Usually, these take the form of *Neumann* boundary conditions, where surface normals are prescribed at perhaps one or two (singular) points inside the domain, as well as at points around the periphery. An *occluding boundary* is particularly useful in this context. Suppose that light exitant from a smooth rock is projected orthographically onto the image. Then, because the object disappears smoothly from view, we may compute normals at those points on the object that correspond to points on the boundary of its image. Given such a boundary point, then the corresponding normal on the object is constrained to be perpendicular to the boundary contour, and to lie in a plane parallel to the image-plane. It is essential that any computational shape-from-shading technique be able to incorporate bounding normals of this kind.

Finally, there are many potential applications of work in this area. These, for example, include automated interpretation of medical, satellite, and synthetic-aperture radar imagery, as well as the enhancement of visual capabilities in robotics. Of course, a further benefit of research in this area may well be an enhanced understanding of the human visual system's processing of shading information.

2. Foundational shape from shading techniques

We now review some of the major iterative techniques employed in shading analysis. It should be noted that the names given to the various methods are not standard in the field. Note additionally that we do not concern ourselves with *local* techniques which employ rather drastic assumptions about surface shape in order to transform the problem from one of global character to one of local character.

2.1 Characteristic strip method

The first automated scheme for solving the shape from shading problem was developed at MIT by Horn [7]. This we now review, using a modified form of Woodham's account [16]. Given the image irradiance equation

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

where $p = u_x$ and $q = u_y$, we may obtain, by differentiation, the pair of equations

$$E_x = p_x R_p + q_x R_q$$
$$E_y = p_y R_p + q_y R_q.$$

Here, subscript again denotes partial differentiation with respect to the given variable. Using the chain rule, small movement in the image may be related to change in surface gradient by the approximations

$$\Delta p \approx \Delta x\, p_x + \Delta y\, p_y$$
$$\Delta q \approx \Delta x\, q_x + \Delta y\, q_y.$$

These equations may be conveniently rewritten in the form

$$\begin{aligned} (\Delta p\,, \Delta q)^T &\approx \mathcal{H}(\Delta x\,, \Delta y)^T \\ (E_x\,, E_y)^T &= \mathcal{H}(R_p\,, R_q)^T, \end{aligned} \tag{1}$$

where $\mathcal{H}$ is the Hessian matrix

$$\mathcal{H} = \begin{pmatrix} p_x & p_y \\ q_x & q_y \end{pmatrix}.$$

The recovery of shape is now based on the following steps:

(i) Assume that R and E are given (although either may be in discrete form), and that at some point (x, y) the values $p(x, y)$ and $q(x, y)$ are known.

(ii) Calculate $R_p\big(p(x,y), q(x,y)\big)$, $R_q\big(p(x,y), q(x,y)\big)$, $E_x(x, y)$ and $E_y(x, y)$, using first-difference approximations, if necessary.

(iii) Choose a constant α such that movement in the image defined by $(\Delta x, \Delta y)$ is small when defined according to the equation

$$(\Delta x\,, \Delta y)^T = \alpha(R_p\,, R_q)^T. \tag{2}$$

The right-hand-sides of equations (1) are thus related by a constant of proportionality.

(iv) Conclude that $(\Delta p, \Delta q)^T = \alpha(E_x\,, E_y)^T$, and that the gradient value at $(x + \Delta x, y + \Delta y)$ is therefore $(p + \Delta p, q + \Delta q)$. Here, the assumption is made that the Hessian remains almost constant over small movement in the xy-plane, and so the left-hand-sides of equations (1) are also related by the constant of proportionality α.

Expressing this formally, we may infer from (2) that

$$(\Delta x^2 + \Delta y^2)^{1/2} = \alpha({R_p}^2 + {R_q}^2)^{1/2}.$$

Let $\tau = (\Delta x^2 + \Delta y^2)^{1/2}$ be a fixed, desirable step size in the image. Then, at any given point (x, y) in the image with gradient $\big(p(x,y), q(x,y)\big)$,

$$\alpha = \tau({R_p}^2 + {R_q}^2)^{-1/2}.$$

We thus arrive at the following iterative method

$$\begin{aligned} \alpha^k &= \tau({R_p}^2 + {R_q}^2)^{-1/2} \\ (x^{k+1}, y^{k+1}) &= (x^k + \alpha^k R_p,\ y^k + \alpha^k R_q) \\ p(x^{k+1}, y^{k+1}) &= p(x^k, y^k) + \alpha^k E_x(x^k, y^k) \\ q(x^{k+1}, y^{k+1}) &= q(x^k, y^k) + \alpha^k E_y(x^k, y^k), \end{aligned}$$

where, again, it should be noted that R_p and R_q are each dependent on the gradient $\big(p(x^k, y^k), q(x^k, y^k)\big)$. This scheme essentially computes normals along base characteristic curves in the image. In this version, successive step sizes in the image are each guaranteed to be of length τ. Prior to computation, values for p and q are required along an initial curve that cuts across the base curves.

In the presence of accurate data, Horn showed that the method performs well. However, drawbacks with the approach are that:

(a) the scheme is susceptible to noise in that errors are likely to be compounded as computation proceeds.

(b) given the necessity of following base characteristic curves, coverage of the domain tends to be uneven, resulting in an inhomogeneous density of normals determined.

(c) the method is not amenable to parallelism on a grid.

(d) coverage of the image domain will only be extensive if the initial curve of starting conditions is appropriately shaped.

2.2 Gradient regularization method

This technique is similar to various early attempts at formulating an iterative scheme suitable for parallel implementation on a grid. It is considered here because it captures some essential features and pitfalls of this approach, and serves to motivate successive approaches.

We seek a solution, defined over the image domain Ω, to the equation

$$E(x,y) - R\big(p(x,y), q(x,y)\big) = 0.$$

This we may express in variational terms as the need to find functions p and q, defined over Ω, that minimise the functional

$$\iint_{\Omega} \Big(E(x,y) - R\big(p(x,y), q(x,y)\big)\Big)^2 \, dx\, dy.$$

This expression may be regarded as a means of measuring the difference between the actual image, and the image implied by a given choice of functions p and q. However, the minimisation problem as stated may not be well-posed (in the sense of having a unique solution shape): in general, many function pairs might minimise the integral. Accordingly, we adopt a regularization term [13] that essentially increases the likelihood of the minimisation problem being well-posed. the new component of the functional penalises change in p and q, leading to the following:

$$\iint_{\Omega} \Big(E(x,y) - R\big(p(x,y), q(x,y)\big)\Big)^2 + \lambda\big(p_x^2 + p_y^2 + q_x^2 + q_y^2\big) \, dx\, dy.$$

The two error measures are balanced by the choice of the scalar λ. The Euler equations for this minimisation problem can be determined and then simplified to read

$$(E - R)R_p + \lambda \nabla^2 p = 0$$
$$(E - R)R_q + \lambda \nabla^2 q = 0,$$

where

$$\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$$

is the Laplacian operator. A solution to the Euler equation is assumed to minimise the functional. Accordingly, we attempt to solve the Euler equations by employing first-difference approximations to yield the iterative scheme

$$p_{i,j}^{k+1} = \overline{p}_{i,j}^{k} + \frac{\epsilon^2}{4\lambda}\big(E_{i,j} - R(p_{i,j}^{k}, q_{i,j}^{k})\big)R_p(p_{i,j}^{k}, q_{i,j}^{k})$$
$$q_{i,j}^{k+1} = \overline{q}_{i,j}^{k} + \frac{\epsilon^2}{4\lambda}\big(E_{i,j} - R(p_{i,j}^{k}, q_{i,j}^{k})\big)R_q(p_{i,j}^{k}, q_{i,j}^{k}).$$

Here, ϵ is the spacing between picture cells, and the approximation

$$\{\nabla^2 p\}_{i,j} \approx \frac{4}{\epsilon^2}(\overline{p}_{i,j} - p_{i,j})$$

is adopted, where the local average, $\overline{p}_{i,j}$, is given by

$$\overline{p}_{i,j} = \tfrac{1}{4}(p_{i,j+1} + p_{i+1,j} + p_{i,j-1} + p_{i-1,j}).$$

Ikeuchi & Horn have observed [11] that schemes of this type exhibit improved stability and convergence when the smoothed values $\overline{p}_{i,j}$ and $\overline{q}_{i,j}$ are used instead of $p_{i,j}$ and $q_{i,j}$ as parameters to R and its various derivative functions.

This approach to the shading problem has the obvious merit of being suited to parallel implementation on a grid. It will also operate in the presence of noisy data, given the minimisation formulation, enabling a "best" solution to be determined. There are, however, some serious problems with the scheme presented:

(a) No account is taken of the fact that the sought functions p and q are dependent. Thus, it is likely that a pq–field converged upon will fail to correspond to a genuine surface.

(b) Ideally, a functional should be designed so that solutions evaluate to zero. Failing this, the minimisation procedure may well go "beyond" the correct solution towards a functional value of zero. In the case that functions p and q are found that cause the above functional to evaluate to zero, then $E(x,y) - R\big(p(x,y), q(x,y)\big) = 0$ and $\big(p_x^2 + p_y^2 + q_x^2 + q_y^2\big) = 0$ everywhere in the domain. The former equality is, of course, desirable. However, the latter will only be satisfied by a plane! Thus, even when given the correct solution as its starting shape, the scheme will tend to move away from (or "warp") it somewhat.

(c) The scheme is not appropriate for use when occluding boundaries arise, since at least one of the gradient functions will increase without limit as the boundary is approached.

(d) There is no proof of convergence for the scheme.

(e) Minimisation approaches such as this are susceptible to local minima, given that the search space is unlikely to be wholly convex.

3. Accommodating the occluding boundary

Having previously noted the importance of the occluding boundary, we now look at two methods that, unlike the previous technique, are able to handle bounding normals that are in the plane of the image.

3.1 Stereographic parameter method

This approach, due to Ikeuchi & Horn [11], was the first to employ explicitly the variational calculus. It is closely related to the gradient regularization technique (although preceeding it), but differs in that an alternative parameterisation is employed.

If we consider a unit sphere resting on the origin of an xy-plane, then surface normals on the southern hemisphere may be mapped onto the plane by gnomonic projection. Here, a line is drawn from the centre of the sphere to a point in the plane. This point then represents the normal that is situated at the intersection of line and sphere. The values of x and y in such a mapping correspond directly to the gradient values of the associated normal; that is, $p = x$ and $q = y$. It may be seen from this visualisation that, as we approach the equator of the sphere, at least one of p and q increases without limit. Thus, a drawback of the gradient parameterisation approach is its inability to incorporate critical boundary conditions associated with the occluding boundary discussed earlier.

If, instead of the gnomonic projection, we employ stereographic projection, this time drawing a line from the north-pole of the sphere to the plane below, we uniquely associate normals on the southern hemisphere with points (f, g) in a circle of radius 2 centred on the origin. Our two parameterisations are then related by the equations

$$f = \frac{2p}{1 + \sqrt{1 + p^2 + q^2}} \qquad \text{and} \qquad g = \frac{2q}{1 + \sqrt{1 + p^2 + q^2}}.$$

We may now develop a technique employing f and g. Let the reflectance map now be parameterised on f and g. For convenience, the name R is retained despite its altered dependency. We use the functional

$$\iint_{\Omega} \Big(E(x,y) - R\big(f(x,y), g(x,y)\big)\Big)^2 + \lambda \left(f_x^2 + f_y^2 + g_x^2 + g_y^2\right) dx\, dy$$

which is now to be minimised with respect to f and g. The analysis now proceeds analogously to that presented in the section on the gradient regularization method, resulting in the scheme

$$f_{i,j}^{k+1} = \overline{f}_{i,j}^{k} + \frac{\epsilon^2}{4\lambda}\big(E_{i,j} - R(f_{i,j}^k, g_{i,j}^k)\big) R_f(f_{i,j}^k, g_{i,j}^k)$$

$$g_{i,j}^{k+1} = \overline{g}_{i,j}^{k} + \frac{\epsilon^2}{4\lambda}\big(E_{i,j} - R(f_{i,j}^k, g_{i,j}^k)\big) R_g(f_{i,j}^k, g_{i,j}^k).$$

Here, as before, ϵ denotes the spacing between picture cells, while $\overline{f}$ and $\overline{g}$ are the local averages of f and g.

This approach overcomes the problem of dealing with the occluding boundary. However, other drawbacks given in the gradient regularisation method are still present.

3.2 Unit-normal method

In this approach, Brooks & Horn [3] adopt a unit-normal parameterisation. The reflectance map R is now assumed to be configured with respect to the unit-normal function $\mathbf{n}(x,y)$. We seek to minimise the functional

$$\iint_{\Omega} \Big(E(x,y) - R\big(\mathbf{n}(x,y)\big)\Big)^2 + \lambda\big(\|\mathbf{n}_x\|^2 + \|\mathbf{n}_y\|^2\big) + \mu(x,y)\big(\|\mathbf{n}\|^2 - 1\big)\,dx\,dy,$$

with respect to $\mathbf{n}(x,y)$. Here, $\mu(x,y)$ is a Lagrangian multiplier function used to impose the constraint that $\mathbf{n}(x,y)$ be a unit vector. The associated Euler equation reduces to

$$(E - R)R_{\mathbf{n}} + \lambda\nabla^2\mathbf{n} - \mu\,\mathbf{n} = \mathbf{0},$$

from which we derive the scheme

$$\begin{aligned}\mathbf{m}_{i,j}^{k+1} &= \overline{\mathbf{n}}_{i,j}^{k} + \frac{\epsilon^2}{4\lambda}\big(E_{i,j} - R(\mathbf{n}_{i,j})\big)R_{\mathbf{n}}(\mathbf{n}_{i,j}) \\ \mathbf{n}_{i,j}^{k+1} &= \mathbf{m}_{i,j}^{k+1} \Big/ \big\|\mathbf{m}_{i,j}^{k+1}\big\|.\end{aligned}$$

As usual, we have that

$$\bar{\mathbf{n}}_{i,j} = \tfrac{1}{4}(\mathbf{n}_{i,j+1} + \mathbf{n}_{i,j-1} + \mathbf{n}_{i+1,j} + \mathbf{n}_{i-1,j}).$$

In the specific case that a Lambertian surface is illuminated by a point source in direction $\mathbf{s}$, we obtain the reflectance map

$$R_{\mathbf{s}}\big(\mathbf{n}(x,y)\big) = \mathbf{n}(x,y)\cdot\mathbf{s}.$$

Substituting this directly into the aforementioned result, we obtain in this special situation the scheme

$$\begin{aligned}\mathbf{m}_{i,j}^{k+1} &= \overline{\mathbf{n}}_{i,j}^{k} + \frac{\epsilon^2}{4\lambda}(E_{i,j} - \mathbf{n}_{i,j}^{k}\cdot\mathbf{s})\,\mathbf{s} \\ \mathbf{n}_{i,j}^{k+1} &= \mathbf{m}_{i,j}^{k+1} \Big/ \big\|\mathbf{m}_{i,j}^{k+1}\big\|.\end{aligned}$$

Here, the mode of operation of the scheme is apparent: for each point, a new normal is computed by taking an average of normals in the neighbourhood, and adjusting this either towards or away from the light source by an amount dependent on the brightness error of the estimate.

Brooks and Horn also developed a method that is aimed at recovering both shape and source direction from shading information. The functional

$$\iint_\Omega \Big(\big(E - \mathbf{n}\cdot\mathbf{s}\big)^2 + \lambda(\|\mathbf{n}_x\|^2 + \|\mathbf{n}_y\|^2) + \mu(x,y)(\|\mathbf{n}\|^2 - 1) \Big)\, dx\, dy$$

is now minimised not only with respect to the function $\mathbf{n}(x,y)$, but also with respect to the variable $\mathbf{s}$. Details of the latter minimisation are omitted here (see [3]), but the resulting coupled scheme is

$$\begin{aligned}
\mathbf{m}_{i,j}^{k+1} &= \overline{\mathbf{n}}_{i,j}^{k} + \frac{\epsilon^2}{4\lambda}(E_{i,j} - \mathbf{n}_{i,j}^{k}\cdot\mathbf{s}^k)\,\mathbf{s}^k \\
\mathbf{n}_{i,j}^{k+1} &= \mathbf{m}_{i,j}^{k+1} / \left\|\mathbf{m}_{i,j}^{k+1}\right\| \\
\mathbf{s}^{k+1} &= \Big(\sum_{i,j\in\Omega} \mathbf{n}_{i,j}^{k+1}\mathbf{n}_{i,j}^{k+1\,T} \Big)^{-1} \sum_{i,j\in\Omega} E_{i,j}\,\mathbf{n}_{i,j}^{k+1}.
\end{aligned}$$

This performs reasonably well in simple situations, is elegantly formulated, and overcomes the occluding boundary problem. It requires less information about the scene conditions than many other techniques. However, other drawbacks previously mentioned still apply. It should be noted that the coupled scheme faces increased prospect of succumbing to local minima in the search space.

4. Incorporating integrability

The following techniques seek to ensure that the generated vector field of normals corresponds to a genuine surface.

4.1 Gradient integrability method

This method reverts to the problematic gradient parameterisation, but compensates by tackling the problem of the non-integrability of gradient fields discussed previously.

In the gradient regularisation method, we do not take into account the fact that, for C^2 surfaces, the gradient functions p and q are related over Ω by the equation

$$p_y(x,y) - q_x(x,y) = 0.$$

In order to remedy this, Horn & Brooks [9] employ the functional

$$\iint_\Omega \big(E(x,y) - R(p,q)\big)^2 + \lambda\big(p_y - q_x\big)^2\, dx\, dy,$$

in which departure from integrability is penalised. The corresponding Euler equations may now be derived and written as

$$\begin{aligned}
(E-R)R_p + \lambda(p_{yy} - q_{xy}) &= 0 \\
(E-R)R_q + \lambda(q_{xx} - p_{yx}) &= 0.
\end{aligned}$$

Employing the usual finite-difference approximations, we arrive at the scheme

$$p_{i,j}^{k+1} = \hat{p}_{i,j}^{k} - \tfrac{1}{2}\widetilde{q}_{i,j}^{k} + \frac{\epsilon^2}{2\lambda}\big(E_{i,j} - R(p_{i,j}^{k}, q_{i,j}^{k})\big)R_p(p_{i,j}^{k}, q_{i,j}^{k})$$
$$q_{i,j}^{k+1} = \hat{q}_{i,j}^{k} - \tfrac{1}{2}\widetilde{p}_{i,j}^{k} + \frac{\epsilon^2}{2\lambda}\big(E_{i,j} - R(p_{i,j}^{k}, q_{i,j}^{k})\big)R_q(p_{i,j}^{k}, q_{i,j}^{k}),$$

where

$$\hat{p}_{i,j} = \tfrac{1}{2}(p_{i,j+1} + p_{i,j-1}) \qquad \text{and} \qquad \hat{q}_{i,j} = \tfrac{1}{2}(q_{i+1,j} + q_{i-1,j})$$

are the vertical average of p and the horizontal average of q, respectively, while $\widetilde{p}_{i,j}$ and $\widetilde{q}_{i,j}$ are estimates of the cross derivatives (times ϵ^2) obtained using the approximations

$$\widetilde{p}_{i,j} = \tfrac{1}{4}(p_{i+1,j+1} + p_{i-1,j-1} - p_{i-1,j+1} - p_{i+1,j-1})$$
$$\widetilde{q}_{i,j} = \tfrac{1}{4}(q_{i+1,j+1} + q_{i-1,j-1} - q_{i-1,j+1} - q_{i+1,j-1}).$$

It has been noted in [9] that this scheme is almost equivalent to that of Strat [15]. The latter method was the first successful parallel scheme on a grid and was derived using a conventional (as opposed to a variational) minimisation technique.

We observe the following:

(a) When the correct solution is given as initial values for $p_{i,j}$ and $q_{i,j}$, the scheme has the obviously desirable property of warping the surface hardly at all. This is in marked contrast to the previously listed variational schemes which are often prone to flattening correct initial surfaces.

(b) The use of an integrability penalty term by no means ensures that the resulting field will be integrable. If the minimum value of the functional is non-zero, perhaps due to a noisy image, then the integrability relation may not hold everywhere. Clearly, the penalty function does not in general enforce integrability.

(c) Convergence of the scheme is much slower than the aforementioned methods.

(d) The inability to handle the occluding boundary is apparent once again.

A detailed discussion of these points may be found in [9].

4.2 Fourier projection method

A quite different approach to the devising of an integrable iterative scheme is taken by Frankot & Chellappa [5]. Here a correction mechanism is employed that enforces integrability. In particular, the method cycles through the following steps:

(i) Use any given shading scheme to generate a new normal field.

(ii) If necessary, convert the normals to pq-form. The resulting field $\big(p(x,y), q(x,y)\big)$ will not, in general, correspond to a gradient field (such that $p_y = q_x$).

(iii) Project the non-integrable field (p, q) onto the nearest gradient field $(\mathbf{p}, \mathbf{q})$.

(iv) Convert the new gradient field back to the parameterization required by the shading algorithm (assuming the scheme is not already in pq).

Clearly, the novel step here is (iii), and this we now describe in brief. Assume that we may usefully represent a surface $z(x, y)$ by the finite expansion

$$z(x,y) = \sum_{\omega \in \Omega} C(\omega)\, e^{i\left(\omega \cdot (x,y)\right)}$$

where $\omega = (\omega_1, \omega_2)$ is a two-dimensional index, Ω is a finite set of indices across the image domain, $e^{i(\omega_1 x + \omega_2 y)}$ are Fourier, orthonormal basis functions, and $C(\omega)$ are Fourier coefficients. It then follows that the gradient may be expressed as:

$$z_x(x,y) = \mathbf{p}(x,y) = \sum_{\omega \in \Omega} i\omega_1 C(\omega)\, e^{i(\omega_1 x + \omega_2 y)}$$
$$z_y(x,y) = \mathbf{q}(x,y) = \sum_{\omega \in \Omega} i\omega_2 C(\omega)\, e^{i(\omega_1 x + \omega_2 y)}$$

Now, suppose that p and q are non-integrable estimates of the gradient values. These may be represented in the form

$$p(x,y) = \sum_{\omega \in \Omega} i\omega_1 C_1(\omega)\, e^{i(\omega_1 x + \omega_2 y)}$$
$$q(x,y) = \sum_{\omega \in \Omega} i\omega_2 C_2(\omega)\, e^{i(\omega_1 x + \omega_2 y)}$$

Note that, in this case, our non-integrable estimates do not have the same Fourier coefficients. We now seek the gradient field $(\mathbf{p}, \mathbf{q})$ that is nearest to our non-integrable field (p, q). This we may do by minimising the functional

$$\iint_{\Omega} |\mathbf{p} - p|^2 + |\mathbf{q} - q|^2 \, dx\, dy.$$

with respect to $\mathbf{p}$ and $\mathbf{q}$. Substituting into this measure the expressions for $\mathbf{p}$, $\mathbf{q}$, p, and q, and minimising with respect to the coefficients C, it may be shown that

$$C(\omega) = \frac{\omega_1^2 C_1(\omega) + \omega_2^2 C_2(\omega)}{\omega_1^2 + \omega_2^2}.$$

We are therefore in a position to determine the nearest integrable field by substituting this expression into the above equations for $\mathbf{p}$ and $\mathbf{q}$..

Problems with this approach are that the gradient parameterisation is retained, and that the Fourier technique requires periodic boundary conditions. The latter factor constitutes a severe limitation. Furthermore, there is no a priori reason to suppose that this projection method will result in final estimates of improved quality.

4.3 Height projection method

This method may be regarded as a variation of the previous technique, except that projection is now achieved by finding an actual surface that best matches the current estimates of the normals.

Given a possibly non-integrable field (p, q), we may determine the graph z that best fits this field by minimising the following functional

$$\iint_\Omega (z_x - p)^2 + (z_y - q)^2 \, dx \, dy$$

with respect to z. This leads to the Euler equation

$$\nabla^2 z = p_x + q_y,$$

from which the iterative scheme

$$z_{i,j}^{k+1} = \overline{z}_{i,j}^{k} - \frac{\epsilon}{8}(p_{i+1,j} - p_{i-1,j} + q_{i,j+1} - q_{i,j-1})$$

is readily generated. Thus, the "nearest" surface to the field (p, q) is calculated, and from this an integrable field $(\mathbf{p}, \mathbf{q})$ may now be derived by taking the appropriate first-differences of z. This method was advanced by Simchony et al. [14].

4.4 Height and gradient method

In this integrated approach, not unrelated to the height projection method, Horn [8] and (independently) Zheng & Chellappa [17] seek to minimise the functional

$$\iint_\Omega \Big(E(x,y) - R\big(p(x,y),q(x,y)\big)\Big)^2 + \lambda\big(p_x^2 + p_y^2 + q_x^2 + q_y^2\big) \\ \mu\big((z_x - p)^2 + (z_y - q)^2\big) \, dx \, dy$$

with respect to z, p and q. The associated Euler equations may be simplified to read

$$\begin{aligned} \lambda\nabla^2 p + (E - R)R_p + \mu(z_x - p) &= 0 \\ \lambda\nabla^2 q + (E - R)R_q + \mu(z_y - q) &= 0 \\ \nabla^2 z - p_x - q_y &= 0. \end{aligned}$$

A discrete approximation of these equations then results in the scheme

$$\begin{aligned} p_{i,j}^{k+1} &= \alpha\overline{p}_{i,j}^{k} + \beta(z_{i+1,j} - z_{i-1,j}) + \gamma\big(E_{i,j} - R(p_{i,j}^k, q_{i,j}^k)\big)R_p(p_{i,j}^k, q_{i,j}^k) \\ q_{i,j}^{k+1} &= \alpha\overline{q}_{i,j}^{k} + \beta(z_{i,j+1} - z_{i,j-1}) + \gamma\big(E_{i,j} - R(p_{i,j}^k, q_{i,j}^k)\big)R_q(p_{i,j}^k, q_{i,j}^k) \\ z_{i,j}^{k+1} &= \overline{z}_{i,j}^{k} - \epsilon/8(p_{i+1,j} - p_{i-1,j} + q_{i,j+1} - q_{i,j-1}) \end{aligned}$$

where $\alpha = \theta/(\theta + \mu)$, $\beta = \mu/\big(2\epsilon(\mu + \theta)\big)$, $\gamma = 1/(\theta + \mu)$, $\theta = 4\lambda/\epsilon^2$, and $\overline{p}_{i,j}$, $\overline{q}_{i,j}$, and $\overline{z}_{i,j}$ are the usual local averages.

An important idea associated with this scheme is that as a solution is approached, the parameter λ is made to fall away to zero. In this way, the regularisation factor gives the early benefit of a rapid "push" towards a solution without the later disregard for integrability. See Harris [6] for related work concerning surface interpolation.

5. Open problems

5.1 Convergence

An important property of an iterative scheme is that it be provably convergent. A further requirement is, of course, that the estimate converged upon be accurate: a convergent scheme is of little value if it settles upon an estimate that is far from the correct solution.

Proof of the convergence of an iterative scheme in this field is rare, as is indicated by the dearth of papers in the area (however, see [12, 4]). A major problem here is that the commonly adopted functionals are not readily amenable to the standard numerical-analysis methods used to prove convergence.

A similar problem arises in relation to the need for a scheme to generate a good estimate. Here it is necessary to avoid converging to estimates that "reside" in local minima. This can be done by designing a functional that is convex in nature, with the desired solution at the (single) minimum point in the space. However, it is very difficult to achieve this within the constraints of the problem. For example, the brightness-error measure

$$\iint_{\Omega} \left(E(x,y) - R(p(,y),q(x,y))\right)^2$$

seems to be a natural and essential component, yet it is almost certain to lead to non-convexity of the overall functional (depending on the choice of R). There are in addition many subtleties associated with the significance of Euler equations (see, for example, Courant & Hilbert [?]).

Another problem arises in relation to methods that reduce the effect of regularisation over time. It is unclear whether the remaining portion of the iterative scheme has "push" in the direction of the solution. Further investigation into the numerical aspects of shading schemes is clearly needed.

5.2 Boundary conditions

Given a particular shape-from-shading problem, a reasonable question arises as to whether it is well-posed (assuming that we are provided, say, with an image, some knowledge of the scene conditions, and prescription of some of the solution surface's normals). Here we think of well-posed as meaning the problem has a unique solution shape. Since recovery of shape from shading can, with sufficient information, be formulated as a first-order partial differential equation, the problem is thus equivalent to that of determining whether the partial differential equation has a unique

solution, given the data and boundary conditions. This is a notoriously difficult mathematical problem on which progress has been made only for the relatively simple situation in which a point source illuminates a Lambertian surface, since this then reduces to an eikonal equation. (However, results here are nonetheless very valuable.) More work needs to be done in this area of uniqueness and ambiguity so as to generalise the few results already obtained (see [1] and its references; see also [2] for consideration of the related existence problem).

5.3 Treatment of occluding boundary and integrability

An unresolved problem in shading analysis has been the design of a scheme that is sensitive both to the problems of integrability and the occluding boundary. Recall that a field of 3-vectors parameterised on gradient should have the property that $p_y = q_x$ if it is to correspond to the projection of a C^2 surface's normals. It is therefore appropriate that this equation should, in some form or another, influence the design of an iterative scheme. However, the pq-parameterisation does not permit occluding boundary normals to be incorporated. This suggests that an alternative formulation of integrability should be considered. Horn & Brooks [9] investigated just this problem, and their results are now described in brief.

Recall that earlier we used stereographic parameters in the derivation of a shading scheme. These parameters are related to gradient values according to the equations

$$p = \frac{4f}{4 - f^2 - g^2} \qquad \text{and} \qquad q = \frac{4g}{4 - f^2 - g^2},$$

so that $p_y - q_x = 0$ may be expressed as

$$\frac{f_y(4 + f^2 - g^2) - g_x(4 - f^2 + g^2) + 2(g_y - f_x)fg}{(4 - f^2 - g^2)^2} = 0.$$

This equation, even when multiplied by $(4 - f^2 - g^2)^2$, is rather unwieldy, leading to very complicated Euler equations. These in turn offer little hope of generating tractable iterative schemes.

We may instead seek an integrability constraint using the unit-normal parameterisation. Let $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$ and $\hat{\mathbf{z}}$ denote unit vectors in the x, y and z directions, respectively. We have that

$$p = -\frac{\mathbf{n} \cdot \hat{\mathbf{x}}}{\mathbf{n} \cdot \hat{\mathbf{z}}} \qquad \text{and} \qquad q = -\frac{\mathbf{n} \cdot \hat{\mathbf{y}}}{\mathbf{n} \cdot \hat{\mathbf{z}}},$$

from which we may determine that

$$p_y = -[\mathbf{n}\, \mathbf{n}_y\, \hat{\mathbf{y}}]/(\mathbf{n} \cdot \hat{\mathbf{z}})^2$$
$$q_x = +[\mathbf{n}\, \mathbf{n}_x\, \hat{\mathbf{x}}]/(\mathbf{n} \cdot \hat{\mathbf{z}})^2.$$

(Here $[\mathbf{a}\, \mathbf{b}\, \mathbf{c}]$ denotes the vector triple product $\mathbf{a} \cdot (\mathbf{b} \times \mathbf{c})$, see [9] for detailed steps.) The constraint $(p_y - q_x) = 0$ can thus be written as

$$\frac{1}{(\mathbf{n} \cdot \hat{\mathbf{z}})^2}\big([\mathbf{n}\, \mathbf{n}_x\, \hat{\mathbf{x}}] + [\mathbf{n}\, \mathbf{n}_y\, \hat{\mathbf{y}}]\big) = 0.$$

A useful "non-integrability" measure to employ in a functional might therefore be

$$\left([\mathbf{n}\,\mathbf{n}_x\,\hat{\mathbf{x}}] + [\mathbf{n}\,\mathbf{n}_y\,\hat{\mathbf{y}}]\right)^2.$$

This is pursued in [9] with a resulting iterative scheme that is unwieldy and awkward.

The challenge therefore remains to generate an iterative scheme sensitive both to the occluding boundary and the issue of integrability.

Acknowledgements

The comments of A. Branford, D. Gibbins and especially R. Kozera were appreciated.

References

[1] Brooks, M.J., W. Chojnacki & R. Kozera (1991a) "Circularly-Symmetric Eikonal Equations and Non-Uniqueness in Computer Vision," *J. Math. Anal. Appl.*, in press.

[2] Brooks, M.J., W. Chojnacki & R. Kozera (1991b) "Shading Without Shape," *Quarterly of Appl Math.*, in press.

[3] Brooks, M.J. & B.K.P. Horn (1985) "Shape and Source from Shading," *Proceedings of the International Joint Conference on Artificial Intelligence*, Los Angeles, CA, August 18–23, pp. 932–936. (Also appears in [10].)

[4] Durou, J.D & H. Maitre (1990) "On the Convergence of Strat's Shape from Shading Method," Departement Images, Telecom Paris 46, rue Barrault, 75013 Paris.

[5] Frankot, R.T. & R. Chellappa (1988) "A Method for Enforcing Integrability in Shape from Shading Algorithms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 10, No. 4, pp. 439–451, July. (Also appears in [10].)

[6] Harris, J.G. (1987) "A New Approach to Surface Reconstruction: the Coupled Depth-Slope model," *Proc. Int. Conf. Computer Vision*, pp. 277-283, London, June.

[7] Horn, B.K.P. (1975) "Obtaining Shape from Shading Information," Chapter 4 in *The Psychology of Computer Vision*, P.H. Winston (ed.), McGraw-Hill Book Company, New York, NY, pp. 115–155. (Also appears in [10].)

[8] Horn, B.K.P. (1990) "Height and Gradient from Shading," *International Journal of Computer Vision.* Vol. 5, No. 1, pp. 37–75, August.

[9] Horn, B.K.P. & M.J. Brooks (1986) "The Variational Approach to Shape from Shading," *Computer Vision, Graphics and Image Processing*, Vol. 33, No. 2, pp. 174–208, February. (Also appears in [10].)

[10] Horn, B.K.P. & M.J. Brooks (eds.) (1989) *Shape from Shading*, MIT Press, Cambridge, Mass.

[11] Ikeuchi, K. & B.K.P. Horn (1981) "Numerical Shape from Shading and Occluding Boundaries," *Artificial Intelligence*, Vol. 17, No. 1–3, pp. 141–184, August. Also in *Computer Vision*, Brady, J.M. (ed.), North Holland Publishing Company, Amsterdam, Netherlands. (Also appears in [10].)

[12] Lee, D. (1988) "A Provably Convergent Algorithm for Shape from Shading," Image Understanding Workshop, Miami Beach, FL, December 9–10, pp. 489–496. (Also appears in [10].)

[13] Poggio, T.A., V.R. Torre & C. Koch (1985) "Computational Vision and Regularization Theory," *Nature*, Vol. 317, No. 26, pp. 314–319, September.

[14] Simchony, T. & R. Chellappa (1987) "Direct Analytical Methods for Solving Poisson Equations in Computer Vision Problems," *Proceedings of the IEEE Computer Society Workshop on Computer Vision*, Miami Beach, FL, November 30–December 2, pp. 44–50.

[15] Strat, T. (1979) "A Numerical Method for Shape from Shading for a Single Image," S.M. thesis, Department of Electrical Engineering and Computer Science, MIT, Cambridge MA.

[16] Woodham, R.J. (1978) "Reflectance Map Techniques for Analyzing Surface Defects in Metal Castings," Ph.D. Thesis, Department of Electrical Engineering and Computer Science, MIT. Also Technical Report TR-457, Artificial Intelligence Laboratory, MIT, Cambridge, MA 02139, June.

[17] Zheng, Q. & R. Chellappa (1990n) "A Robust Algorithm for Inferring Shape from Shading," USC-SIPI Report No. 159, Signal and Image Processing Institute, Dept of Elec. Eng. Systems, Univ. of Southern California.

Discipline of Computer Science
School of Information Science and Technology
The Flinders University of South Australia
GPO Box 2100, Adelaide, SA 5001, Australia