Acta Applicandae Mathematicae **30**: 125–157, 1993.
© 1993 *Kluwer Academic Publishers. Printed in the Netherlands.*

Some Mathematical Problems in Computer Vision*

A. M. BRUCKSTEIN
Department of Computer Science, Technion – Israel Institute of Technology, 32000 Haifa, Israel

and

A. TANNENBAUM

Department of Electrical Engineering, University of Minnesota, Minneapolis, MN 55455, U.S.A.

(Received: 20 July 1992)

Abstract. Several interesting mathematical problems arising in computer vision are discussed. Computer vision deals with image understanding at various levels. At the low level, it addresses issues like segmentation, edge detection, planar shape recognition and analysis. Classical results on differential invariants associated to planar curves are relevant to planar object recognition under partial occlusion, and recent results concerning the evolution of closed planar shapes under curvature controlled diffusion have found applications in shape decomposition and analysis. At higher levels, computer vision problems deal with attempts to invert imaging projections and shading processes toward depth recovery, spatial shape recognition and motion analysis. In this context, the recovery of depth from shaded images of objects with smooth, diffuse surfaces require the solution of nonlinear partial differential equations. Here results on differential equations, as well as interesting results from low-dimensional topology and differential geometry are the necessary tools of the trade. We are still far from being able to equip our computers with brains capable to analyze and understand the images that can easily be acquired with camera-eyes; however the research effort in this area often calls for both classical and recent mathematical results.

Mathematics Subject Classifications (1991). 68T10, see: 53A04, 35F25, 35Q80, 57N25, 58A99.

Key words. Differential invariants, planar curve evolutions, shape from shading, photometric stereo.

1. Introduction

Recent advances in technology enabled computers to acquire images with inexpensive high resolution cameras, creating a pressing need to interpret and analyze visual data. In the beginning, this need was met by various ad-hoc algorithms, some quite successful. However, quite early on there also was a trend to import advanced mathematical results to attack some of the better understood and defined problems of

*This work was supported in part by NSF grant DMS-8811084, Air Force Office of Scientific Research Grant AFOSR-90-0024, and the Army Research Office DAAL03-91-G-0019, and by the Technion Fund for Promotion of Research.

the field. Today, group theory, discrete geometry, differential and algebraic geometry, topology, and partial differential equations have become the tools of the trade. Indeed, it is generally agreed in the vision community that the import of mathematics to the field of computer vision helps greatly in understanding some of the key problems, and focuses the research on the need to work and experiment within precisely defined and accepted paradigms.

There are several directions in computer vision research. Some very interesting problems arising in visual inspection of industrial products and robotics lead to research in so-called *low level vision.* Work in model based shape analysis and recognition has already resulted in many useful products, such as optical character recognizers, handwriting recognizer interfaces to computers, printed-circuit board inspection systems and quality control devices. In spite of such successes many low-level vision problems remain to be addressed. Efficient ways for analyzing, recognizing, and understanding even planar shapes, when they do not come from a well-defined and documented library of shapes, when they are distorted by a geometric viewing transformation, such as perspective projection or when they are partially occluded, have yet to be developed.

The mainstream of research in computer vision are the so-called *high level* topics, such as depth perception from single or stereo images, the use of shading clues to infer object shapes and surfaces, the use of image sequences to recognize objects in motion and analyze dynamic three-dimensional scenes, with a view to applications like robot navigation and collision avoidance. Such problems require good understanding of the imaging processes, i.e., the projection geometry, light propagation and surface reflectance properties. While many of these problems are mathematically well-defined, they often lead to ill-posed inverse problems. For example, the imaging process produces images of objects, and we want to extract the shape of these three dimensional objects from such images. Certain intriguing facets of questions in this domain can be successfully modeled and dealt with mathematically, however it is fair to state that the problems that can be mathematically adequately treated at this point are special cases, i.e., heavily simplified versions of the real issues to be dealt with. Sometimes the simplifications that lead to well-posed mathematical problems are reasonable in view of controlled environments in robotics problems; good examples are the problems arising from active vision and structured light research. The problems involved in developing general-purpose vision algorithms with human-like performance still seem as formidable as they were when the first computers were built. A very promising trend in addressing this issue is the placement of emphasis on qualitative rather than quantitative image understanding.

In this paper we shall discuss four problems in computer vision, emphasizing the mathematical tools that were found to be useful in dealing with them. The problems are motivated by real issues and their mathematical solution should be viewed as a guidance toward possible real-world implementation and application. The problems are:

(1) recognition of planar objects with smooth boundaries under partial occlusion, when the objects are distorted by general viewing transformations;
(2) planar object analysis and decomposition by boundary evolution;
(3) recovering the shape, or the depth to the surface of a diffuse three-dimensional object from a shaded image taken under controlled illumination conditions;
(4) shape recovery of diffuse objects from two images taken under different illumination conditions (photometric stereo).

The mathematical results that found applications to the above problems are the theory of differential invariants for smooth planar curves, the theory of planar curve evolution under curvature-dependent differential deformation laws, some elements from the theory of nonlinear first-order partial differential equations, and some basic consequences of integrability on smooth surfaces. The four sections that follow will describe these problems and their solutions in some detail.

2. Local Invariants and Planar Curve Recognition

The problem of recognizing and locating a partially visible planar object, whose shape is distorted by a viewing transformation arises in several machine vision tasks. Such problems raise the question of invariants under viewing transformations. We describe here a theory of local invariants of smooth planar curves under projective, affine and similarity transformations. The roots of this theory go back to some classical results of Elie Cartan; see [W], [L], [SB]. We follow the approach described in [BN] to the use of local invariants for recognition of partially occluded planar objects.

2.1. PROBLEM DEFINITION

A simple, closed planar curve, the, say $[0, 1]$, to points in $\mathbf{R}^2$. The curve may also be regarded as the trajectory of a point moving in the plane, the position at time t being given by $P(t) = [x(t), y(t)]$. We have, by assumption, $P(1) = P(0)$ and $P(t_i) \neq P(t_j)$ for any $t_i \neq t_j$; $t_i, t_j \in [0, 1]$. Further assume the boundary curves and their traversal are smooth, implying that the functions $x(t)$, $y(t)$ are differentiable several times. Obviously, a simple closed planar trajectory in the plane may be traversed at various speeds and therefore $P(\tilde{t}(t))$, where $\tilde{t}(t) = t_0 + \phi(t)$ and $\phi(t)$ is a smooth monotone function $\phi\colon [0, 1] \to [0, 1]$, describes the same trajectory as $P(t)$, with a different initial position and different traversal schedule. Such elementary transformations of planar curve descriptions are called curve *reparametrizations*. To separate the geometric concept of a planar curve from its formal algebraic description, we refer to the planar curve described by $P(t)$ as the image of $P(t)$, $\mathrm{Im}\{P(t)\}$. Denoting $P(\tilde{t}(t))$ by $R_{\tilde{t}(t)}[P(t)]$ we have that

$$\mathrm{Im}\{P(t)\} = \mathrm{Im}\{R_{\tilde{t}(t)}[P(t)]\}, \tag{2.1}$$

for any $\tilde{t}(t) = t_0 + \phi(t)$ as above; thus all smooth reparametrizations are equally good descriptions of any given curve. For a planar curve we may choose to work with any valid traversal as its formal description. Suppose that the points of $\mathbf{R}^2$ are subjected to a geometric transformation, $T_\psi: \mathbf{R}^2 \to \mathbf{R}^2$,

$$T_\psi[(x, y)] = [X_\psi(x, y), Y_\psi(x, y)] \tag{2.2}$$

parametrized by a vector of parameters ψ. A planar curve will be distorted by T_ψ, and the points of $\mathrm{Im}\{P(t)\}$ will be mapped to another simple and closed curve in the plane. Choosing an arbitrary parametric description for the distorted curve, $\tilde{P}(\tilde{t})$, we have that

$$\begin{aligned}\tilde{P}(\tilde{t}) &= [\tilde{x}(\tilde{t}), \tilde{y}(\tilde{t})] = T_\psi[x(\tilde{t}), y(\tilde{t})] \\ &= T_\psi[R_{\tilde{t}(t)}[x(t), y(t)]] = T_\psi[R_{\tilde{t}(t)}[P(t)]] = R_{\tilde{t}(t)}[T_\psi[P(t)]],\end{aligned} \tag{2.3}$$

i.e. the curve description $\tilde{P}(\tilde{t})$ is always a $T_\psi[\cdot]$-distorted version of a reparametrization of $P(t)$, or equivalently, a reparametrization of a $T_\psi[\cdot]$-distorted version of $P(t)$. We shall analyze ways to account for the consequences of looking at a planar object with smooth boundaries from various unknown points of view. This induces several types of geometric transformations, $T_\psi[\cdot]$, that distort the boundaries. The questions we address are:

(1) Given a library of planar objects and a distorted view of one of them, recognize (identify) the object from its distorted image (see Figure 1a).
(2) Given a library of objects and the profile of a cluster of objects from the library, distorted by possibly different viewing transformations, resolve the cluster into its components (see Figure 1b).

In both problems above, we assume the distortions to be of a given class $T_\psi[\cdot]$, with no knowledge of its parameters ψ. The most general geometric transformations on planar shapes that we shall deal with are the so-called *projective mappings*. They arise in the context of the laws of perspective projection, and can be best described by representing points in the plane in homogeneous coordinates, as follows

$$[x, y] \to [x\lambda^{-1}, y\lambda^{-1}, \lambda^{-1}] \tag{2.4}$$

where the third coordinate is an arbitrary scaling factor. Using homogeneous coordinates, a planar object or curve is lifted into projective space. To $P(t) = [x(t), y(t)]$ we may associate a projective curve described by $[x(t)\lambda^{-1}(t), y(t)\lambda^{-1}(t), \lambda^{-1}(t)]$, where $\lambda(t)$ is any continuous smooth function with $\lambda(t) > 0$. The general projective transformation is then any linear mapping applied to the "lifted" trajectory, i.e.

$$[X(t), Y(t), Z(t)] = [x(t)\lambda^{-1}(t), y(t)\lambda^{-1}(t), \lambda^{-1}(t)]\mathbf{A} \tag{2.5}$$

where $\mathbf{A} = [a_{ij}]$ is any full rank matrix. Notice that the arbitrary scaling function chosen, $\lambda(t)$, multiplies all entries of $[X(t), Y(t), Z(t)]$, and projecting this homogenized curve back to its 2D representation we obtain

$$\tilde{P}(t) = [\tilde{x}(t), \tilde{y}(t)] = \left[\frac{a_{11}x(t) + a_{21}y(t) + a_{31}}{a_{13}x(t) + a_{23}y(t) + a_{33}}, \frac{a_{12}x(t) + a_{22}y(t) + a_{32}}{a_{13}x(t) + a_{23}y(t) + a_{33}}\right]. \tag{2.6}$$

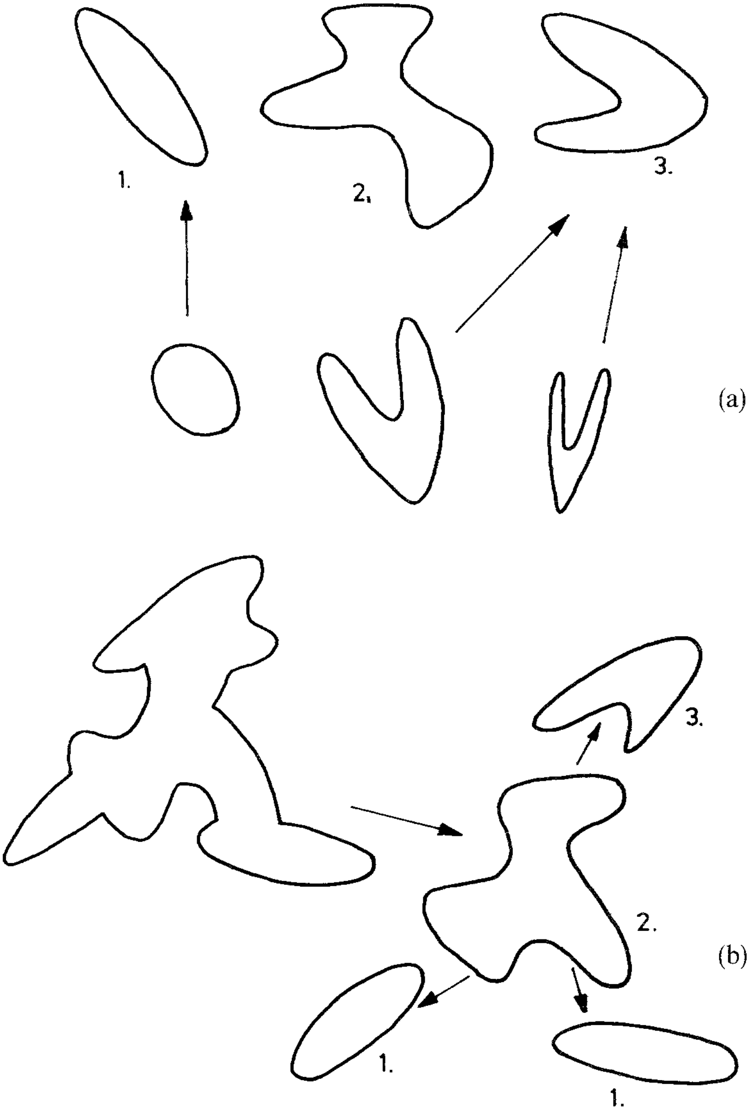

Fig. 1. (a) Shape identification, (b) cluster resolution.

This transformation is the most general one encountered as the embedding of a viewing transformation. The important particular cases of this transformation that are usually analyzed in detail are the rigid motions in the plane, similarity transformations and affine mappings. The equations describing *rigid* (*Euclidean*) *motion* mappings in the plane are

$$[\tilde{x}, \tilde{y}, 1] = [x, y, 1] \begin{bmatrix} \cos\omega & -\sin\omega & 0 \\ \sin\omega & \cos\omega & 0 \\ v_x & v_y & 1 \end{bmatrix}, \tag{2.7}$$

the parameters of the transformation being the rotation angle ω, that defines a rotation matrix $\mathrm{Rot}(\omega)$, and the translation vector $V = [v_x, v_y]$. The equations for *similarity transformations* are the same as the ones for rigid motions with the rotation matrix $\mathrm{Rot}(\omega)$ in (2.4) replaced by $\alpha\,\mathrm{rot}(\omega)$, where α is a scaling factor. The parameters in this case are ω, v_x, v_y and α. *Affine mappings* are defined by a general non singular matrix A, replacing the rotation matrix $\mathrm{Rot}(\omega)$ in the definition of rigid motions. The parameters of an affine transformation are the entries of A, and the translation vector V.

These are the geometric transformations $T_\psi[\cdot]$. The general projective map is defined by 8 parameters (one of the entries of $\mathbf{A}$ may be normalized to 1, with no effect on the 2D $\to$ 2D transformation). This map generalizes both a perspective viewing transformation and an affine map. Note however that a true perspective projection has fewer parameters. The affine map has 6 independent parameters, while the similarity transformation has 4 and rigid (Euclidean) motions are characterized by 3 parameters. Note that for all the above transformations there exists a parameter choice ψ that maps the transformation $T_\psi[\cdot]$ into the identity transformation $I[\cdot]$, i.e. $I[P(t)] = P(t)$. In fact, since the matrices involved are invertible, the transformation types discussed are continuous groups of transformations.

2.2. CANONICAL CURVE PARAMETRIZATIONS VIA DIFFERENTIAL INVARIANTS

Suppose we are given a planar object with a smooth boundary, the image of a closed planar curve described by $P(t)$. If the object is subjected to a geometric transformation of the type discussed in the previous section, the transformed planar object will have boundary that can be described by $\tilde{P}(\tilde{t})$

$$\tilde{P}(\tilde{t}) = T_\psi[R_{\tilde{t}(t)}[P(t)]] \equiv R_{\tilde{t}(t)}[T_\psi[P(t)]]. \tag{2.8}$$

We do not know the parameters ψ of the geometric transformation and we are only given the images of two closed boundary curves in the plane, $\mathrm{Im}\{P(t)\}$ and $\mathrm{Im}\{\tilde{P}(\tilde{t})\}$.

To solve the first object recognition problem posed above, we must be able to decide whether an arbitrary description $\tilde{P}(\tilde{t})$ could be related to $P(t)$ via the equation (2.5), for some reparametrization $\tilde{t}(t)$ and some transformation parameters ψ. In order to solve the second, more difficult, cluster resolution problem we should be able to identify even *portions* of a given curve $\tilde{P}(\tilde{t})$ as transformed and reparameterized portions of an original shape boundary described by $P(t)$.

An example: The Euclidean Invariants

In order to illustrate a general approach to attack these problems, let us first analyze the way they are solved for the simplest case of rigid (Euclidean) motion transformations. It is well-known that a smooth planar curve has an intrinsic curvature versus arc-length representation $k(s)$. The arc-length is a rotation-translation invariant and so is the curvature. Therefore, in so representing a closed contour the only arbitrary choices are an initial position on the curve and the direction of traversal. If there are no unambiguously defined 'landmark' points on the boundary, or in the case of partial occlusion situations, the initial point will remain arbitrary. The direction of traversal may usually be chosen *a priori*. In any case, all shapes that are rotated and translated versions of the original will have boundaries described by $k(s - s_0)$, i.e., translated versions of the same intrinsic-description function. Thus, both the object recognition and the cluster resolution problem may be solved via a total or partial correlation, or 1D function-matching (string matching), process. The curvature versus arc-length representation solves both our problems by first devising a *curve-dependent reparametrization* of the boundary curve and associating to the curve a *signature function* that is invariant under the given class of geometric transformations. The curve-dependent reparametrization, given the description $P(t) = [x(t), y(t)]$, is readily obtained as

$$s(t) = \int_0^t \left[\left(\frac{\mathrm{d}}{\mathrm{d}\xi} x(\xi) \right)^2 + \left(\frac{\mathrm{d}}{\mathrm{d}\xi} y(\xi) \right)^2 \right]^{1/2} \mathrm{d}\xi \tag{2.9}$$

by using the metric

$$\mathrm{d}s = [(\mathrm{d}x(t)/\mathrm{d}t)^2 + (\mathrm{d}y(t)/\mathrm{d}t)^2]^{1/2}\,\mathrm{d}t \tag{2.10}$$

and, after reparametrizing $P(t)$ to get $P(s) = [x(s), y(s)]$, the curvature invariant $k(s)$ is given by

$$k(s) = \frac{\mathrm{d}x(s)}{\mathrm{d}s}\frac{\mathrm{d}^2y(s)}{\mathrm{d}s^2} - \frac{\mathrm{d}^2x(s)}{\mathrm{d}s^2}\frac{\mathrm{d}y(s)}{\mathrm{d}s} := x^{(1)}(s)y^{(2)}(s) - x^{(2)}(s)y^{(1)}(s). \tag{2.11}$$

Having obtained the curvature invariant, we may ask whether we can obtain others, independent of it (since clearly we could obtain other invariants by performing various operations on $k(s)$). Such questions are central to the theory of differential invariants.

The above procedure for identifying planar objects from portions of their boundaries in case of (Euclidean) rigid motion transformations serves as a model for the type of solutions that we seek in the case of general transformations. Given a parametric family of transformations on $\mathbf{R}^2$, $T_\psi[x, y] = [\tilde{x}, \tilde{y}]$, and a planar curve $P(t) = [x(t), y(t)]$, we have to determine a reparametrization,

$$\mathrm{d}\tau = \Gamma\{P(t)\}\,\mathrm{d}t, \tag{2.12}$$

so that for $\tilde{P}(\tilde{t}(t))$ we shall have

$$\mathrm{d}\tilde{\tau} = \Gamma\{\tilde{P}(\tilde{t})\}\,\mathrm{d}\tilde{t}\,|_{\tilde{t}} = \mathrm{d}\tau. \tag{2.13}$$

Here $\Gamma\{P(t)\}$ is a positive function of $x(t)$, $y(t)$ and their derivatives, $x^{(k)}(t)$, $y^{(k)}(t)$, $(k = 1, 2, 3, \dots)$, i.e., it depends on the local behavior of the curve at the point $P(t)$. Reparametrizing both $P(t)$ and $\tilde{P}(\tilde{t})$ by τ and $\tilde{\tau}$, respectively, we have from (2.13) that

$$\tilde{P}(\tilde{\tau}) = T_\psi[P(\tilde{\tau} + \tau_0)] \tag{2.14}$$

and the next step is the search for a signature function invariant under $T_\psi[\cdot]$. Suppose we find a transformation Λ mapping $P(\tau)$ into a function $\rho(\tau)$, based also on the local behavior of the curve $P(\tau)$

$$\rho(\tau) = \Lambda[P(\tau)], \tag{2.15}$$

so that the function $\rho(\cdot)$ is an invariant signature function, i.e.,

$$\rho(\tau) = \Lambda[P(\tau)] = \Lambda[\tilde{P}(\tilde{\tau}(\tau))] = \tilde{\rho}(\tau - \tau_0). \tag{2.16}$$

Then, given a curve $P(t)$ that undergoes a T_ψ transformation and a reparametrization $\tilde{t}(t)$, to yield $\tilde{P}(\tilde{t}) = T_\psi[R_{\tilde{t}(t)}[P(t)]]$, we may define the function $\rho(\tau)$ associated to $P(t)$ to be a generalized 'curvature' versus 'arc-length' representation of $P(t)$. We first need to find a function of the local behavior of the planar curve that transforms, under reparametrization, as follows

$$\Gamma\{P(t)\} = \Gamma\{\tilde{P}(\tilde{t}(t))\} \frac{\mathrm{d}\tilde{t}(t)}{\mathrm{d}t}. \tag{2.17}$$

Then with the reparametrizations (2.12) and (2.13) we shall have that $\tilde{\tau} = \tau_0 + \tau$, and for this reparametrization, we have by (2.17)

$$\Gamma\{P(\tau)\} = \Gamma\{\tilde{P}(\tilde{\tau})\} \cdot 1. \tag{2.18}$$

It might seem that we have also produced an invariant signature function. However, applying (2.17) to the identity transformation $T_\psi[\cdot] = I[\cdot]$ we obtain

$$\Gamma\{P(\tau)\} = \Gamma\{P(t)\} \cdot \frac{\mathrm{d}t}{\mathrm{d}\tau} = \Gamma\{P(t(\tau))\} \frac{1}{\Gamma\{P(t(\tau))\}} = 1, \tag{2.19}$$

showing that we have associated a trivially invariant signature (a constant) to the curve, that hardly was worth working for. If, however, we could obtain two *different* functions Γ_1 and Γ_2 both satisfying (2.17), then we could use one of them for reparametrization and the second for an invariant signature, since then

$$\Gamma_2\{P(\tau)\} = \Gamma_2\{P(t)\} \frac{\mathrm{d}t}{\mathrm{d}\tau} = \frac{\Gamma_2\{P(t)\}}{\Gamma_1\{P(t)\}}, \tag{2.20}$$

and using (2.17) repeatedly we have

$$\Gamma_2\{\tilde{P}(\tilde{\tau})\} = \Gamma_2\{\tilde{P}(\tilde{t})\} \frac{\mathrm{d}\tilde{t}}{\mathrm{d}\tilde{\tau}} = \frac{\Gamma_2\{\tilde{P}(\tilde{t})\}}{\Gamma_1\{\tilde{P}(\tilde{t})\}} = \frac{\Gamma_2\{P(t(\tilde{t}))\}}{\Gamma_1\{P(t(\tilde{t}))\}} = \Gamma_2\{P(\tilde{\tau} - \tau_0)\}. \tag{2.21}$$

This is a key observation showing how to use differential invariants for shape recognition. Suppose we managed to provide for a class of geometric transformations a generalized 'ρ-curvature' vs. 'τ-arc-length' representation. This representation enables us to locate corresponding points on the curves $\mathrm{Im}\{P(t)\}$ and $\mathrm{Im}\{\tilde{P}(\tilde{t})\}$, i.e., if we are given some point Ω on $\mathrm{Im}\{P(t)\}, \Omega = P(t_\Omega)$, we can look for $\tilde{\Omega} = T_\psi\{\Omega\}$ by locating the point on $\tilde{P}(\tilde{t})$ that the same value for the generalized curvature $\rho(\tau)$. Locating the corresponding points for several Ω_i on $\mathrm{Im}\{P(t)\}$ and writing the equations

$$\tilde{\Omega}_i = T_\psi\{\Omega_i\}, \tag{2.22}$$

we obtain a system of equations for the parameters of the geometric transformations, and in some cases we might be able to uniquely determine ψ. Therefore, we could use two images of a planar object to identify the geometric transformation that affected an image under consideration. This could be done even if the boundary of the object is only partially visible in the image distorted by a viewing transformation, since the above-discussed generalized curvature versus arc-length representations are based on the local behavior of the boundaries.

Several papers in the computer vision literature deal with problems of object recognition under distorting, geometric viewing transformations. According to whether it was assumed that the entire object is visible in the distorted image, or only portions of it, we may classify the approaches to such problems as based on global or local information. When global information is available we could attempt to identify the parameters ψ of the geometric transformation by analyzing how global shape parameters, like perimeters, areas, higher order moments, etc., are affected by a T_ψ transformation. For object recognition we may also rely on so-called global invariants associated to shapes; these are quantities that remain invariant when the planar shape undergoes a T_ψ transformation. The search for global invariants under various geometric transformations is an ongoing concern of current pattern recognition research. However, here we discussed methods that employ only local information. It is only through such methods that one can solve object identification problems under partial occlusion. An approach to using local information that is very popular in the pattern recognition literature advocates the use of special-feature points on the boundary, such as breakpoints, ends of straight portions or inflection points. Such points can relatively easily be located based on local information and can also be *identified* on transformed object boundaries. Then they may serve either for segmentation or for the identication of the transformation T_ψ. These methods however, are unsuitable for smooth boundaries and in cases where occlusion wipes out feature points. It is then worthwhile to study generalized, invariant signatures or, ρ-curvature versus τ-arc-length type representations. The pair of functions $\Gamma_1\{\cdot\}$ and $\Gamma_2\{\cdot\}$, are provided for the common viewing transformations by a theory of differential invariants, developed in the late 19th and early 20th centuries by Halphen (see [L]), Wilczynski [W], and Cartan (see [SB]). For more details and references see [IW1], [COCD], [BN].

Unfortunately, from a practical point of view, the problem with this approach is that high derivatives of the curve representations are required to produce the $\Gamma\{\cdot\}$ functions (the highest derivative needed being by one less than the number of parameters of the continuous group of deformations). Thus there is a need to move away from differential to 'neighborhood behavior'-based invariants and here the classical mathematical theories no longer seem to help. Thus, we must find ways to exploit the local properties of the curve to yield invariant functionals to be used for recognition. As an example see the paper [BKLP], for tricks enabling similarity invariant recognition of partially occluded boundaries, and [VKO] for a Lie-group approach to semi-differential invariants for recognition, reducing the number of derivatives needed by exploiting some point-match information that might be available. For further developments see also [BPB], [IW2] and [BHNR].

3. Shape Analysis via Boundary Evolution Equations

When we look at a planar shape we usually give it an interpretation, such as a conglomerate of some elementary parts or as a basic shape with some structured protuberances. Very early on in computer vision, researchers wanted to understand and model mathematically the way in which we make such 'gestalt' decisions, that seem to be crucial in the process of associating meaning to the picture we see. Planar shapes are described in the computer in various ways, using either contour or region descriptions, each of them making different types of information explicit. While complete descriptions are obviously equivalent we would like to have algorithms to choose a description that is readily suited for the task at hand and algorithms that use the description of choice to come up with the 'right' interpretations. Blum [HB] argued that a useful concept for shape analysis is the notion of the 'shape skeleton', a planar graph, made of curve portions retaining some of the characteristics of a shape, and perhaps enabling its decomposition into meaningful parts. The *skeleton* of a planar shape is formally defined as the set of points, whose minimal distance to the boundary is attained at two or more different points on it. This concept is perhaps best understood via a dynamic, so-called 'prairie fire model' [HB], [RCC]: assume that the shape under consideration is set on fire simultaneously along its boundary and the fire propagates in the normal direction at unit speed. The skeleton is the set of points where the fire quenches (Figure 2), and it can be shown that complete reconstruction of the original shape is possible given the skeleton and the so-called 'quenching function', providing the distance from each skeleton point to the outside world! Below we show the differential equation describing such a prairie fire propagation. It is easy to see that the skeleton is defined by the shock fronts that this equation generates. An extension of such a shape analysis process was found to be very useful in providing reasonable interpretations to planar shapes. We shall outline this approach for representing and analyzing shape, an approach based on a parametrized class of boundary evolution equations. It is based on the recent work of [BK], [KTZ1], [KTZ2]. One of the main ideas in this work is to show that some natural requirements on boundary evolution are realized by a modification of the prairie fire model. The evolution equations postulated enable an explicit treatment of singularities, based on a series of conservation laws. As the boundaries evolve, singularities develop and separate the shape into subregions, a natural hierarchy of 'parts' emerges. This approach, which is inherently non-linear, stands in contrast and generalizes some earlier linear (e.g., heat equation) scale-space models, incorporating direct handling of discontinuous events.

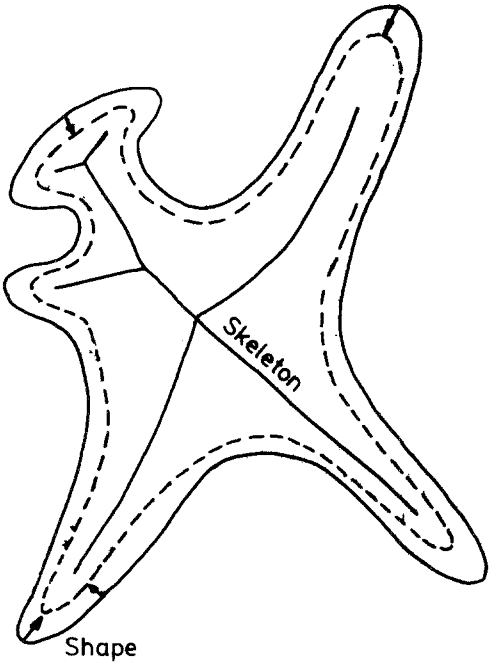

Fig. 2. 'Prairie fire' propagation yielding the skeleton.

The differential equations which we postulate for boundary evolution were previously used in mathematical physics as models for the evolution of phase boundaries [Gur] (and references therein), as well as models of curve shortening [Gr], [GH]. The numerical properties of implementing such equations were thoroughly analyzed by Sethian and Osher [Se1], [Se2], [OS]. The subject has much mathematical depth and is rapidly developing, and we will touch on a number of the more relevant aspects to vision here.

3.1. CURVATURE-DEPENDENT EVOLUTION EQUATIONS FOR PLANAR CURVES

We will be considering families of closed embedded curves $C: S^1 \times [0, t_f) \to \mathbf{R}^2$ evolving according to functions of the curvature. (Since this discussion is informal, we will be rather loose with the hypotheses we will be putting on the families we consider here. See [An] and [KTZ1] for the formal mathematical treatments.) The general deformation of a curve in the plane of interest in vision problems may be given by

$$\frac{\partial C}{\partial t} = \alpha(s, t)T + \beta(s, t)N, \tag{3.1}$$

where N is the (outward) normal, T is tangent, and α, β are smooth functions. Notice that since we are only interested in the images of curves, we may take $\alpha = 0$. (Changing

α leads to a curve reparametrization, without affecting its shape.) Furthermore, we shall constrain the deformations to be determined by the local geometry of the curve, i.e., $\beta(s, t) = G(\kappa)$ where $\kappa(s, t)$ denotes the (Gaussian) curvature of the curve $s \to C(s, t)$. Thus we are led to the following equation:

$$\frac{\partial C}{\partial t} = G(\kappa)N. \tag{3.2}$$

In the mathematics literature, a number of cases for the function G have been explored. For example, there has been a great deal of work in connection with the geometric heat equation in which $G(\kappa) = -\kappa$. In this case, the isoperimetric ratio L^2/A of the curve approaches 4π as the enclosed area approaches 0, and thus the curve shrinks and becomes a circle while disappearing; see [GH], [Gr]. The function of interest to us is

$$G(\kappa) = 1 - a\kappa, \tag{3.3}$$

where $a \geqslant 0$. The case $a = 0$, i.e., $G \equiv 1$ is very important, and here Equation (3.2) becomes an equation that has been studied in relation to problems in geometric optics [Ar1], [Ar2], flame propagation [Se1], and shape morphology [HB], as well as shape decomposition. Indeed, this is the differential equation for the prairie fire model described above. The κ part gives a diffusive effect, while the constant part gives a wave (hyperbolic) effect which tends to create singularities and break a shape into its constituent parts.

Let us first discuss some general properties of Equation (3.2). We first introduce some standard notation. Let

$$\rho(s, t) := \left|\frac{\partial C}{\partial t}\right| = [x_s^2 + y_s^2]^{1/2}, \tag{3.4a}$$

denote the distance measure along the curve. The arc-length parameter s is then

$$\mathrm{s}(s, t) = \int_0^s \rho(\zeta, t)\,\mathrm{d}\zeta. \tag{3.4b}$$

Let the positive orientation of a curve be defined so that the interior is to the left when traversing the curve. The *tangent*, *curvature*, *normal*, *orientation* and *length* are defined in the standard way. We will take the normal to be pointing outward, where the inward or outward is determined by the interior, or equivalently by the orientation of the curve. We then have that

$$T := \frac{\partial C}{\partial \mathrm{s}} = \frac{1}{\rho}\,\partial\,\frac{C}{\partial}\,s, \tag{3.5a}$$

$$\kappa := \frac{1}{\rho}\left|\frac{\partial T}{\partial s}\right|, \tag{3.5b}$$

$$N := \frac{-1}{\kappa\rho}\frac{\partial T}{\partial s}, \tag{3.5c}$$

$$L(t) := \int_0^{2\pi} \rho(s, t)\,\mathrm{d}s. \tag{3.5d}$$

Finally, we let

$$\bar{\kappa}(t) := \int_0^{2\pi} |\kappa(s, t)|\,\rho(s, t)\,\mathrm{d}s \tag{3.6}$$

denote the total absolute Gaussian curvature.

The behavior of the classical of solutions Equation (3.2) can be rather thoroughly analyzed, and one can prove useful results (from the applied point of view) of the following type [An], [KTZ1], [KTZ2]:

1. Let $C(s, t)$ be a classical solution of (1) for $t \in [0, t')$ and $\kappa G(\kappa) \leqslant M$ for all $\kappa \in \mathbf{R}$ (regarding G as a function of κ). Then,

$$L(t) \leqslant \min(L(0) + 2\pi t,\ L(0)\,\mathrm{e}^{Mt}). \tag{3.7a}$$

In particular, for $G(\kappa) = \alpha\kappa - 1$,

$$L(t) \leqslant \min(L(0) + 2\pi t,\ L(0)\,\mathrm{e}^{t/4\alpha}). \tag{3.7b}$$

2. Let $C(s, t)$ be a classical solution of (1) for $t \in [0, t')$. Suppose that $\kappa G(\kappa) \leqslant M$, and $G_\kappa \leqslant 0$. Then

$$\bar{\kappa}(t) \leqslant \bar{\kappa}(0). \tag{3.8}$$

Moreover, if $[0, t')$ is an interval on which a classical solution exists, then one may also show that

$$d_{\mathrm{H}}(C_t, C_0) \leqslant \sqrt{(4at)}, \tag{3.9}$$

where d_{H} denotes the Hausdorff metric on compact subsets of $\mathbf{R}^2$. The above facts imply that

$$\lim_{t \to t'} C_t = C^* \tag{3.10}$$

in the Hausdorff metric, and the curve C^* regarded as a mapping $C^* \to \mathbf{R}^2$ is Holder continuous with exponent 1/2. In the shape-analysis application a major concern are the weak solutions of equations of the type (3.2) which we will now describe.

3.2. CONSERVATION LAWS, SHOCKS, AND ENTROPY CONDITIONS FOR SHAPE ANALYSIS

For the purposes of this survey, let us concentrate on the interesting special case of $G \equiv 1$, i.e. the prairie fire model. Here in the classical manner we will derive the equation describing a *hyperbolic conservation law*. With $G \equiv 1$, we can write Equation (3.2) explicitly as

$$x_t = \frac{y_s}{(x_s^2 + y_s^2)^{1/2}}, \tag{3.11a}$$

$$y_t = -\frac{x_s}{(x_s^2 + y_s^2)^{1/2}}. \tag{3.11b}$$

Note that we are writing $C_t(s) := C(s, t) = [x(s, t), y(s, t)]$ in terms of the position coordinates, the initial curve being $C_0(s) := C(s, 0)$ for $0 \leqslant s \leqslant S$.

Now let us see how from Equations (3.11) we can derive a hyperbolic conservation law. As long as C_t stays smooth and non-self-intersecting, by virtue of the implicit function theorem, we can express the propagating front in the form

$$y = U(t, x). \tag{3.12}$$

(C_t is the graph of (3.12).) Then one can verify that U satisfies the following (Hamilton–Jacobi) equation

$$\frac{\partial U}{\partial t} - (1 + U_x^2)^{1/2} = 0. \tag{3.13}$$

Set

$$u := \frac{\partial U}{\partial x}. \tag{3.14}$$

Then differentiating (3.13) with respect to x we see that

$$u_t - ((1 + u^2)^{1/2})_x = 0. \tag{3.15}$$

Equation (3.15) is in the standard form of a 'hyperbolic conservation law' which has a huge classical and modern literature devoted to it; see [Sm] and the references therein. We shall not go into these laws in depth here, but would like to give the reader some of the physical motivation behind such PDEs.

Explicitly, a *hyperbolic conservation law* is given by the hyperbolic PDE (i.e., a wave-type equation)

$$u_t + F(u)_x = 0. \tag{3.16}$$

The conservation law that (3.16) is expressing mathematically may be formulated as follows: Material is distributed along a line with coordinate x and assume the distribution satisfies the physical conservation law that the temporal rate of change of the amount of material in a fixed interval equals the flux of the material through the boundaries. If u then denotes the density, and $F(u)$ the flux, then mathematically this conservation law may be expressed as

$$\frac{\mathrm{d}}{\mathrm{d}t} \int_x^{\Delta x} u \,\mathrm{d}x = F(u(x)) - F(u(x + \Delta x)). \tag{3.17}$$

Letting $\Delta x \to 0$ gives (3.16), the required conservation law.

Using this interpretation for the moving front given by (3.11), we see that it is the *slope* that is conserved! We should also note that such hyperbolic conservation laws appear in gas dynamics related to the Riemann problem [Sm].

A very important property of such hyperbolic conservation laws (from the point of view of the shape anslysis problem we are studying, and other applications [Sm]) is the fact that one may get discontinuous solutions for (3.11) and (3.6), even in the presence of smooth initial data. Once such discontinuities or *shocks* develop, one must be careful in defining precisely what is meant by a 'solution' to (3.16). There is a notion of *weak* solution for (3.16) by which the discontinuous solution satisfies the PDE (3.6) in a certain distributional sense; see [Sm] for the precise definition. It is easy to see how discontinuities develop for the system (3.11). Indeed, assuming that C_t remains smooth it is easy to recognize that the Gaussian curvature satisfies the following evolution equation:

$$\kappa_t = -\kappa^2. \tag{3.18}$$

We can explicitly solve (3.18) to find that

$$\kappa(s, t) = \frac{\kappa(s, 0)}{1 + t\kappa(s, 0)}.$$

Notice then that if C_0, the initial curve, is anywhere concave, i.e. has Gaussian curvature negative at any point, $\kappa(s, t)$ will blow up in finite time, and the resulting curve will develop corners, i.e., we get a shock.

There is another fundamental problem with the development of shocks, namely that the solution will not, in general, be unique. In order to account for this, one needs a way of picking out the 'physically relevant solution. For example, in the prairie fire case, shocks appear when the wavefront reaches points of quenching. From those points it is quite clear how the wavefronts will propagate, according to the Huygens principle. However mathematically it is not obvious what conditions must be imposed, in order to have the 'correct' weak solutions. The condition ensuring that, at a shock, the weak solution propagated will be the physically meaningful one is the so-called entropy condition, imposing that entropy increases across the shock. The idea is that 'entropy' is the inverse of 'information', and we require that information be lost across the shock. Thus we get a condition of "irreversibility' with the formation of shocks.

In his PhD thesis [Se2], Sethian shows that the condition: *once a particle burns, it remains burnt* (in the prairie fire model), is indeed equivalent to a classical entropy condition (see also [Lax1], [Lax2], [Sm]). This condition prevents a nonself-intersecting curve from evolving into a self-intersecting one. It is a very important fact that such entropy (weak) solutions may also be obtained as limits of (classical) solutions of the hyperbolic Equation (3.16) perturbed by certain viscosity terms. A similar effect ('artificial viscosity') occurs in the discretization of the system. This means that essentially (when everything converges) we pick up such an entropy condition in a correct digital implementation of the hyperbolic conservation law.

It is precisely the formation singularities that allow us to decompose a given figure into its parts. The various possibilities are very nicely analyzed in terms of a classification of the shocks; see [BK]. Suppose we are given a planar shape, and consider it to be the starting point for evolution equations of the type (3.2) the local control function being (3.3), parametrized by a. We obtain a two-dimensional expansion of the given shape into a class of shapes C_t^a, that result from the evolution of the initial shape according to (3.2) with a given for some time t (see Figure 3). The space of evolved shapes together with the singularities encountered and their classification, could conceptually be regarded as a multifaceted, hierarchical description of the original shape, capturing its salient features, and separating the shape into physically meaningful components. One can even imagine ways of defining shape metrics based on the class of evolved shapes, metrics that would be physically and visually natural and meaningful. This research, reported in [BK], [KTZ1], [KTZ2], is still going on. Problems remain to be addressed, such as achieving invariance under affine/viewing transformations, dealing with partially occluded shapes, dealing with projections of genuine three-dimensional shapes and many others. So far, this research produced some very interesting mathematical results and shows much promise for

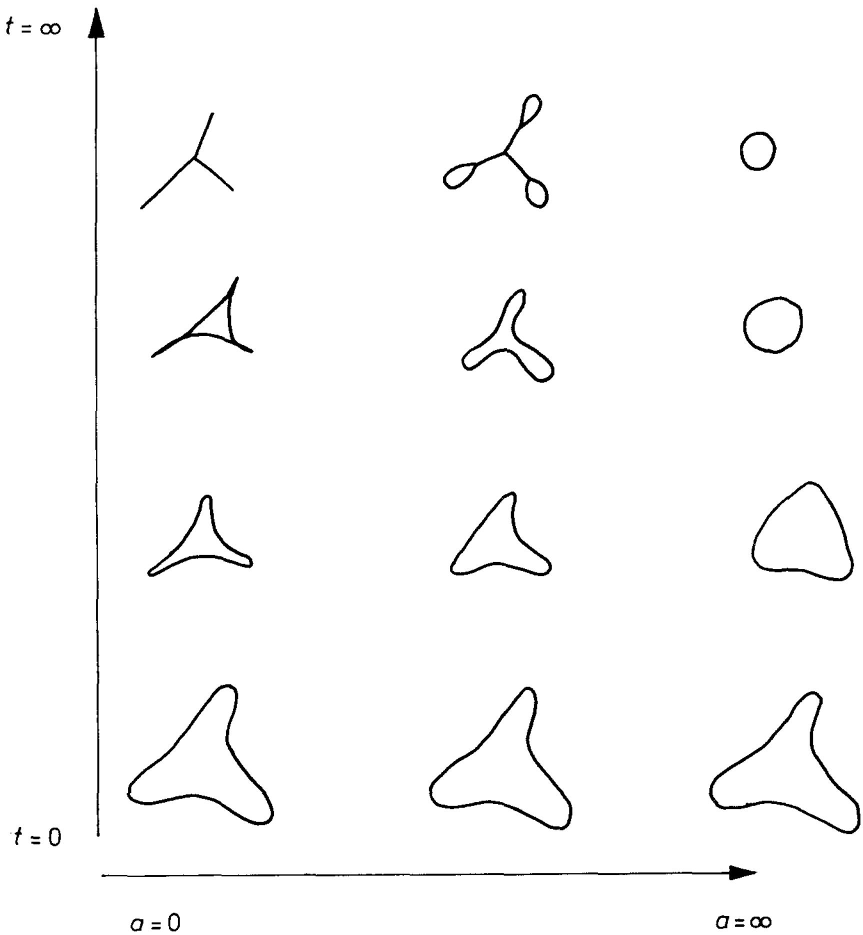

Fig. 3. The shape class corresponding to $\partial C/\partial t = (1 - ak)N$.

planar shape analysis. For some results on affine invariant evolutions and some applications of this theory to Computer Aided Design and Mathematical Morphology see [ST], [BSS], [SKSB] and [KB1].

4. Shape from Shading and the Solution of Nonlinear PDEs

We have seen some of the properties and uses of a planar curve propagation equation of the form

$$\frac{\partial C}{\partial t} = \beta(s, t)N$$

which we shall call the *normal deformation equation* for planar shape analysis. Here we shall see that an entirely different problem in computer vision leads to the same type of equation. However, instead of having a β that depends on the local curvature, we have that it is controlled by a given 'external' function. In this case, from the usual data which arise in vision problems, there will be no shocks, by definition! The problem to which we are referring is the so-called *shape-from-shading* problem. It requires a method for the recovery of a bivariate function $H(x, y)$ that describes a 'nice', almost everywhere differentiable height profile, from a well-defined shading information. The given shading data is assumed to be a result of diffuse, Lambertian reflection of light from the surface. This implies that, if the scene is uniformly illuminated from above, the shading yields information on the cosine of the angle between the vertical and the surface normal at each point. Given the shading information in the plane, the problem is to determine all height profiles consistent with the data, and some boundary conditions, such as points of known height and surface orientation, or height profiles along continuous curves in the image plane. The shape-from-shading method that we discuss is based on a recursive way of determining equal-height or level contours of the surface starting at a given level curve, first presented in [B].

Shape-from-shading problems have a long and interesting history. The first researchers to address the problem of determining shape from shading information were apparently those concerned with the photometric analysis of the lunar topography (see [H1] and the references therein). It is clear that the shading information plays, along with stereo vision and motion clues, an important role in the depth perception process. The theoretical question of how much depth information can be obtained from a single view of a scene from shading alone, thus arises quite naturally. Algorithms, based on the classical characteristic strip method, for determining the depth-profile from a single image produced by various shading rules were considered by B. K. P. Horn and his coworkers in the 1970's [H1-3]. More recent work on shape determination from a single image concentrated on the importance of singular points, surface models and occluding boundaries, in providing initialization and ensuring unique depth recovery. Iterative, relaxation-type techniques were also invented, relying on surface smoothness constraints [Bro], [BKPHB], [IH].

In this section we briefly discuss the basic shape-from-shading problem, and present a method for solving the nonlinear eiconal differential equation involved via evolution equations for *equal-height-contours*. When one such contour is available we can devise a simple algorithm that reconstructs all the equal-height-curves of the surface of interest in a well defined region and also clearly displays the inherent ambiguities of the given problem. This algorithm is easily derived and, in contrast to the classical characteristic-strip expansion method, does not use the derivatives of the shading data. This is a result of a natural way of exploiting lateral constraints in the parallel the propagation of the recovery algorithm. We discuss the ambiguities and possible ways to exploit topological constraints on the behavior of nice surfaces to help the shape recovery process in ambiguous situations.

4.1. THE BASIC SHAPE-FROM-SHADING PROBLEM

Suppose we are given a continuous bivariate function $H(x, y)$, describing a surface in three dimensions, as follows

$$z = H(x, y). \tag{4.1}$$

The shaded image of the surface is $I(x, y)$, the value $I(x, y)$ depending on surface reflectivity, its orientation at (x, y) and illumination conditions. The shape from shading problem that we address is to recover the function $H(x, y)$ over a region D, from the image $I(x, y)$ over that region and possible some further information, e.g., the values of $H(x, y)$ over some continuous curve in D.

The function $I(x, y)$ is defined via a shading rule. It is customary to define the shading rule via a so-called *reflectivity function*, that characterizes the surface properties and provides an explicit connection between $I(x, y)$ and the surface orientation. In the case of a surface with so-called Lambertian diffuse reflection properties and uniform illumination, $I(x, y)$ is simply the cosine of the angle $\alpha(x, y)$, between the surface normal at (x, y) and the direction from which the light falls on the scene. For simplicity we shall always assume that the illumination is uniform and falls on the surface vertically from above.

Define the directional derivatives of the height profile $H(x, y)$ along the x and y directions as

$$p(x, y) = \frac{\partial}{\partial x} H(x, y), \qquad q(z, y) = \frac{\partial}{\partial y} H(x, y). \tag{4.2}$$

The surface normal at (x, y) is clearly perpendicular to the plane determined by the vectors $[1, 0, p]$ and $[0, 1, q]$; therefore it points along the direction of their vector product $[-p, -q, 1]$. The normal vector at (x, y) is thus

$$N(x, y) = \frac{1}{\sqrt{1 + p^2 + q^2}} [-p, -q, 1] \tag{4.3}$$

and the cosine of the angle between $N(x, y)$ and the vertical direction $[0\ 0\ 1]$ is

$$\cos \alpha(x, y) = \frac{1}{\sqrt{1 + p^2 + q^2}}. \tag{4.4}$$

In the Lambertian case with light falling straight from above, we therefore have the shading rule

$$I(x, y) = \frac{1}{\sqrt{1 + p^2 + q^2}} = R_L(p, q). \tag{4.5}$$

Note that the reflectivity function R is defined on the (p, q) plane – called the 'gradient space'. A general (not necessarily Lambertian) shading rule is defined via

$$I(x, y) = R(p(x, y), q(x, y)), \tag{4.6}$$

where $R(\cdot, \cdot)$ is a given function. Equation (4.6) is a nonlinear partial differential equations that has to be satisfied by the surface $H(x, y)$. Therefore solving the shape-from-shading problem amounts to solving a nonlinear partial differential equation, and some boundary conditions are necessary.

Given the image $I(x, y)$, it is in general impossible to unambiguously recover the height profile $H(x, y)$. As an immediate example of ambiguity simply consider the function $-H(x, y)$, which under a Lambertian shading rule, maps into the same image as $H(x, y)$. Some further information on the function $H(x, y)$ is therefore needed. This is usually given as some smoothness constraint on the surface defined by $z = H(x, y)$ (such as C^k continuity) and exact or approximate values of $H(x, y)$ at either a discrete set of points $\{(x_i, y_i)\}$, together with the corresponding surface orientations $\{(p_i, q_i)\}$, or on a continuous curve on the (x, y)-plane (boundary conditions). The given boundary conditions and smoothness assumptions are not always enough to remove ambiguities, and it is in fact very difficult to determine, in general situations, sufficient conditions for a unique solution surface.

4.2. SHAPE-FROM-SHADING VIA EQUAL-HEIGHT CONTOURS

In this section, we will use as data and then try to determine the *equal height contours* of the profile $z = H(x, y)$. An equal height contour or a *level-curve* is a continuous curve in the (x, y)-plane on which the function $H(x, y)$ is constant. If $\{x(\theta), y(\theta)\}$ $\theta \in \Theta$ is the parametric representation of the curve we have

$$\frac{d}{d\theta} H(x(\theta), y(\theta)) = 0. \tag{4.7}$$

One might argue that such a contour contains a lot of information and is scarcely available. In fact, in a number of practical situations one is quite naturally able to determine equal height contours, or portions of them. As an example, the shores of a lake in a landscape readily provide a closed equal height contour; so is the case when an island raises from the sea (any shoreline is an equal height curve). Furthermore, in robot vision systems one might be able to provide illumination which actively delineates one or more equal height contours.

In the sequel the assumption will be that we are given an equal height contour which is almost everywhere differentiable. By definition, along such a curve we have zero height gradient, which yields

$$dH = p\,dx + q\,dy = 0. \tag{4.8}$$

Therefore along the given contour $\{x(\theta), y(\theta)\}$ we have determined a relation between the two directional derivatives of the surface p and q. For almost all θ's we have, rewriting (4.8),

$$p(x(\theta), y(\theta)) \frac{d}{d\theta} x(\theta) = -q(x(\theta), y(\theta)) \frac{d}{d\theta} y(\theta). \tag{4.9}$$

Together with the shading information at that point $I(x(\theta), y(\theta))$, this relation determines p and q, up to an inherent sign ambiguity. Indeed the Lambertian shading rule

$$I(x, y) = R_L(p, q) = \frac{1}{\sqrt{1 + p^2 + q^2}} \tag{4.10}$$

yields $p^2 + q^2 = (1 - I^2)/I^2$, which together with $px' + qy' = 0$ provide

$$p = \pm \frac{y'}{(x'^2 + y'^2)^{1/2}} \frac{(1 - I^2)^{1/2}}{I},$$

$$q = \pm \frac{x'}{(x'^2 + y'^2)^{1/2}} \frac{(1 - I^2)^{1/2}}{I} \quad \text{at } [x(\theta), y(\theta)]. \tag{4.11}$$

We get two pairs of solutions, corresponding to a certain (p, q) vector and its negative counterpart. This is expected since, at each point on the equal height curve, the same grey level would be produced by the shading rule if the tangent plane had the direction of maximum ascent given either by ϕ or by $\phi + \pi$. Note also that we could determine the (p, q) pairs up to a similar ambiguity along *any continuous path* on which the height profile is known *a priori*. In case of equal height curves, the direction of the data contour determines the direction of the maximal surface ascent/descent. Suppose we know that the height profile is a mountain rising from the sea. This immediately settles the direction of the steepest ascent as the vector pointing toward the inner region defined by the equal height contour of the shorelines. Using this information we may determine an equal height contour situated a bit above the sea level, and so on we can recursively climb and reconstruct the height profile – provided no 'problems' occur. Problems arise, as we shall see, if the mountain is not a nice and unimodal profile, and we further discuss these issues after a description of the basic profile reconstruction algorithm.

Assume that $[x,(\theta), y(\theta)]$, $\theta \in [0, 1]$ is a closed curve and that, as θ goes from 0 to 1 we trace the curve in the counterclockwise direction. A tangent vector at θ is simply given by $[x'(\theta), y'(\theta)]$; the unit normal to it pointing inside the curve will be

$$\mathbf{n}_\theta = \frac{1}{[x'(\theta)^2 + y'(\theta)^2]^{1/2}} [-y'(\theta), x'(\theta)]. \tag{4.12}$$

From (4.9) it is clear that in the direction n_θ, we have to go a certain distance d_θ, in order to *climb* a given amount ΔH. If ΔH is small, this distance is quite accurately determined by the shading data alone, since, in the Lambertian example, $I(x(\theta), y(\theta))$ yields the cosine of the angle between the surface normal at $(x(\theta), y(\theta))$ and the vertical direction. As the direction of the maximal ascent is known to be (4.12), we have from geometrical considerations

$$d_\theta = \Delta H \frac{I}{\sqrt{1 - I^2}} \tag{4.13}$$

Therefore, given a closed equal-height contour assumed to be at a reference level H_o, a closed contour situated at the level $H_o + \Delta H$ is determined via (see Figure 4a)

$$\begin{aligned} &[x(\theta, \Delta H), y(\theta, \Delta H)] \\ &\quad = [x(\theta), y(\theta)] + d_\theta \mathbf{n}_\theta \\ &\quad = [x(\theta), y(\theta)] + \frac{1}{\sqrt{x'(\theta)^2 + y'(\theta)^2}} \frac{\Delta H \cdot I_\theta}{\sqrt{1 - I_\theta^2}} [-y'(\theta), x'(\theta)]. \end{aligned} \tag{4.14}$$

This derivation leads to a system of first order nonlinear partial differential equations for the functions $x(\theta, h)$ and $y(\theta, h)$ representing 'doubly parametrized' equal height curves in the (x, y) plane. Indeed, if $[x(\theta, h), y(\theta, h)]$ is defined as a contour corresponding to $H = h$, (4.14) is equivalent to the following set of partial differential equations

$$\begin{aligned} &\frac{\partial}{\partial h} \begin{bmatrix} x(\theta, h) \\ y(\theta, h) \end{bmatrix} \\ &= \frac{I(x(\theta, h), y(\theta, h))}{[1 - I^2(x(\theta, h), y(\theta, h))]^{1/2}} (\mathbf{n}_\theta) \\ &= \frac{I(x(\theta, h), y(\theta, h))}{[1 - I^2(x(\theta, h), y(\theta, h))]^{1/2}} \frac{1}{\left[\left(\frac{\partial}{\partial \theta} x(\theta, h)\right)^2 + \left(\frac{\partial}{\partial \theta} y(\theta, h)\right)^2\right]^{1/2}} \begin{bmatrix} -\frac{\partial}{\partial \theta} y(\theta, h) \\ \frac{\partial}{\partial \theta} x(\theta, h) \end{bmatrix} \end{aligned} \tag{4.15}$$

with initial conditions $[x(\theta, 0), y(\theta, 0)] = [x(\theta), y(\theta)]$.

Note that (4.15) is a *nonlinear initial value problem* of the general form discussed in the previous section. It has to be integrated to obtain the equal height curves of the profile that yielded the shading $I(x, y)$. It is implicit in our derivations that the surface is smooth enough to provide almost everywhere differentiable equal-height contours at all heights h. It is also assumed that those contours are 'well-behaved' as for

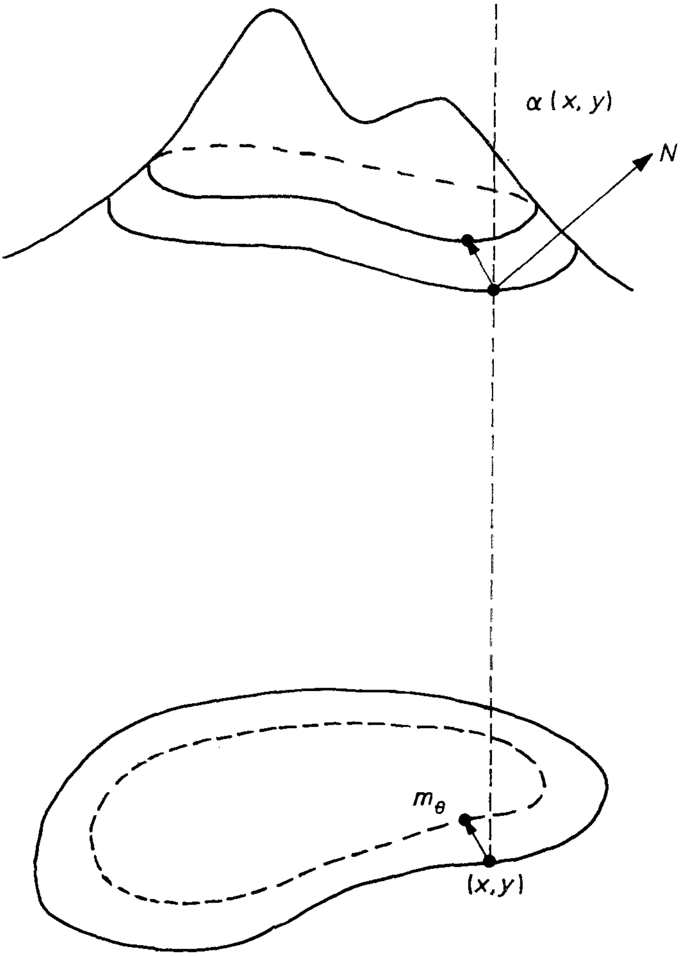

Fig. 4(a). Shape from shading as curve evolution.

example in the case of a unimodal $H(x, y)$ over the region of interest (say the interior of the first equal height contour), when they are nicely nested 'generalized' rings.

It is clear that the recursions (4.14) and their differential counterparts (4.15) are valid generally, provided we are given information on which side of the original equal height curve the surface increases. The data $[x(\theta), y(\theta)]$ can be any curve that is differentiable and if it is not a closed contour we will get, using (4.14), the reconstruction of a well-defined slice of the surface $z = H(x, y)$. If we do start with a closed contour and at some level we obtain a self-intersecting (i.e., not 'well behaved') equal height curve, this means that we encountered a saddle area which separates peaks, or peaks and dips in $H(x, y)$. In this case the contour should be separated into nonintersecting parts and the algorithm may be continued with the separated closed parts as initial equal height curves. An equal height curve may also approach a saddle

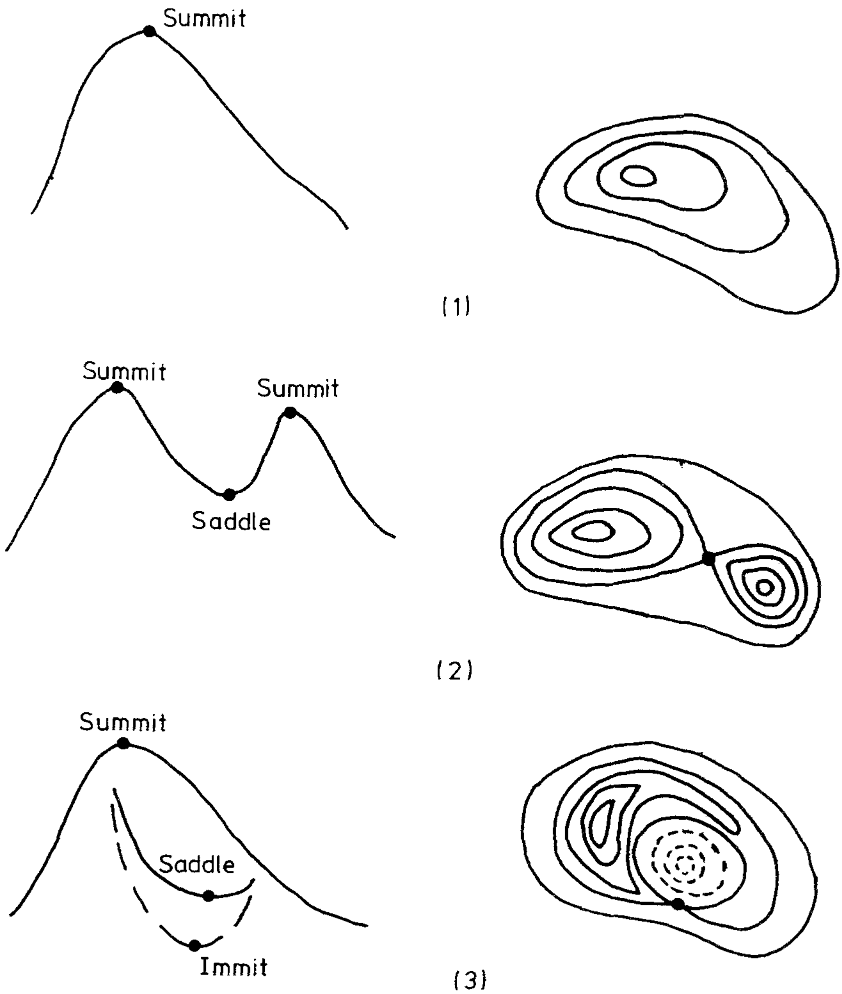

Fig. 4(b). Various types of equal height curve profiles.

point from one side only, and there it will become necessary to continue a partial reconstruction (see Figure 4b).

A thorough discussion of what can happen to the equal-height contour profile of a smooth surface, based on topological constraints can be found in a classical paper by James Clerk Maxwell on 'hills and dales', [Max]. In a modern interpretation, Maxwell proves the so-called 'mountaineers' index theorem, stating that if a surface has isolated simple singularities, i.e., 'summits' (local maxima), 'immits' (local minima) or saddle points, then within an equal-height contour we must have that

$$\#\ \textit{of Summits} + \#\ \textit{of Immits} - \#\ \textit{of Saddle points} = 1$$

The practical implementation of the algorithm will, of course, be based on (4.14), the $[x,(\theta), y(\theta)]$ curve being given (perhaps in a suitably chain-coded way) on a finite grid of θ values. Then we can use several methods to estimate the derivatives $x'(\theta)$ and $y'(\theta)$ that appear in the recursion formula. Also we can leave open the choice of the steps in height (ΔH) taken so as to enable the use of various adaptive schemes (that are particularly useful if the approach to a saddle area is detected). We also note that in practice, $A(x, y)$ is known only on a discrete grid of picture elements or pixels, thus we have to somehow interpolate for the values needed in (4.14) which are points situated on equal height curves. More importantly however, it turns out that the numerical approaches for planar curve evolutions developed by Sethian and Osher, [Se1], [Se2], [OS], are also applicable to the equal height contour evolution based shape from shading, yielding stable and efficient algorithms, see [KB2].

It becomes clear by looking at the system of Equations (4.15) that trouble arises when we approach singular point where $I(x, y) = 1$, indicating $p = q = 0$. A singular point can be either a local extremum or a saddle point of some sort. At an isolated singular point we can define an 'unsafe' neighborhood and when an equal height curve enters such a neighborhood, we disregard that portion of it but continue to propagate the algorithm from the remaining contour. Some portions of the (x, y)-plane will, of course, remain uncovered using this method. The 'singular' areas/curves in the plane for which $I(x, y) = 1$ provide boundaries of possible flips in the directions of maximal ascent and a practical shape from shading process should keep track of these and, based on natural constraints on the behavior of equal-height contours, choose the direction assignments which yield consistent final reconstructions. If *a priori* we know that the surface is unimodal, then no such problems arise, the solution being unique up to height reversal. Note also that we can live with nondifferentiability at a finite set of points along each equal height contour and reconstruct the profile by matching the slices corresponding to differentiable portions. For discussions on uniqueness and existence in shape-from-shading via a dynamical systems approach, see [HBB], [Sax], [O].

5. Two-Image Photometric Stereo and Surface Integrability

Next we shall briefly discuss the problem of recovering the shape of an object with Lambertian surface reflectivity from two images obtained in different illumination conditions, the so-called *photometric stereo problem*. It is easy to see that when two images are available, the local surface normals are ambiguous up to two possible orientations. Following [OB], we shall see that for arbitrary smooth surfaces, the local integrability constraints generically resolve the problem of deciding between the two possibilities. Furthermore, one has a complete characterization of the surfaces that remain ambiguous under the given illumination conditions. As the characterization is dependent on the illumination directions, the two images generically resolve the surface normal ambiguity problem. We shall discuss an algorithm from [OB], for recovering Lambertian surfaces from a pair of images obtained with illumination from two distance sources with known directions. See also [K] and [LB] for some later developments and an interesting application. The reconstruction process has two stages: first the surface normals are recovered, and then the object surface consistent with these normals is reconstructed using a standard depth from normals procedure.

5.1. RECOVERING SURFACE NORMALS FROM PHOTOMETRIC STEREO INFORMATION

The inverse problem of shape-from-shading discussed in the previous section is not well posed, and there might exist a large number of surfaces that could have given rise to a particular image, even under the same conditions of lighting and the same surface reflectance properties. This ambiguity inherent in a single image has been circumvented, by using more or less stringent constraints on the imaged object, or by assuming various types of prior information about it. The classical photometric stereo procedure, see [BKPH], [RW], [M], [I], uses three or more images of an object, taken under different illumination conditions, to locally remove the ambiguities in recovering surface normals. However, as was shown in [OB], for Lambertian surface reflectivity, if only *two* different shaded images of a smooth object are given, its shape can almost always be uniquely determined at all points where no self-shadows occur.

The *two-image* (or *two-source*) photometric stereo problem is the following. We have two images of the same surface, produced without changing the same camera position relative to the surface under different illumination directions. It is required to reconstruct from these images the height profile of the surface.

The model for the generation of image intensities was already mentioned in the previous section. If the height profile is represented by the equation $z = H(x, y)$ and if the function $H(x, y)$ is differentiable, then at each point the normal vector to the surface $N(x, y)$ is given by

$$N(x, y) = \frac{1}{(1 + p^2(x, y) + q^2(x, y))^{1/2}} [-p(x, y), -q(x, y), 1], \tag{5.1}$$

where p and q are the partial derivatives of $H(x, y)$, as in (4.2). The intensity at a point (x, y) in an image of a Lambertian surface, depends, by assumption, only on the angle between the illumination vector and the normal vector at the point. Let A denote the illumination vector, i.e., the unit vector pointing in the direction of the light source. The components of A are a_x, a_y, a_z. Let $\langle A, B\rangle$ denote the scalar product of vectors A and B, and $|B|$ denote the length of a vector B. Then the image intensity I_A at point (x, y) is given by:

$$I_A(x, y) = \langle N(x, y), A\rangle = \frac{1}{(1 + p^2 + q^2)^{1/2}} [(-p(x, y)a_x - q(x, y)a_y + a_z]. \tag{5.2}$$

The second image will be $I_B(x, y)$, described an by expression similar to (5.2), with the components b_x, b_y, b_z of the second illumination vector B, substituted for a_x, a_y, a_z. In order to find the surface orientations these two images will be used. Note that, locally, the partial derivatives must obey the following set of equations:

$$I_A = (-pa_x - qa_y + a_z)(1 + p^2 + q^2)^{-1/2}, \tag{5.3a}$$

$$I_B = (-pb_x - qb_y + b_z)(1 + p^2 + q^2)^{-1/2}. \tag{5.3b}$$

Letting

$$T := (1 + p^2 + q^2)^{1/2} \tag{5.4}$$

and rearranging (5.3) yields

$$pa_x + qa_y = a_z - I_A T, \tag{5.5a}$$

$$pb_x + qb_y = b_z - I_B T. \tag{5.5b}$$

Regarded as two linear equations in two unknowns, these equations can be solved for p and q in terms of T, providing solutions of the form

$$p = c_p T + d_p, \qquad q = c_q T + d_q. \tag{5.6}$$

Recalling the definition of T in (5.4), i.e. $T^2 - p^2 - q^2 = 1$ the solutions for p and q may be inserted providing a quadratic equation for T of the form:

$$K_2 T^2 + K_1 T + K_0 = 0, \tag{5.7}$$

where K_i are functions of I_A, I_B, a_x, a_y, a_z, b_x, b_y and b_z. Solving (5.7) produces two solutions for T, say T_1 and T_2. If the two solutions for T are inserted in (5.6) we obtain two pairs of partial derivatives (p_1, q_1) and (p_2, q_2), corresponding to two normals: N_1 and N_2. This is all that can be obtained using the local constraints provided by two images at the point (x_0, y_0). So far, the only assumption made on the height profile was that it has first derivatives.

The geometric interpretation of the above algebraic manipulations is as follows. The light intensity at each point in an image of a Lambertian surface gives the angle between the normal at that point and the illumination direction. Thus, the locus of all normals that could have produced the intensity I_A at point (x_0, y_0) is a (Monge) cone, with apex at (x_0, y_0) and axis in the illumination direction and having an opening angle determined by $\arccos(I_A)$ (Figure 5a). If the brightness at the same point when illuminated from two different directions (photometric stereo) is known, the normal at (x_0, y_0) must belong to two such cones. Therefore it belongs to their intersection. Two cones with the same apex either intersect along two or one half-lines or do not intersect at all (except for the common apex). The case of no intersection cannot occur for genuine photometric stereo images and will not be considered here. The case of one intersection produces an unambiguous solution, which corresponds to one solution for T in Equation (5.7), and is of some significance as will soon become apparent. The general case is that of two solutions out of which only one is the 'true' normal and this can be seen to agree with the algebra above.

Note that, given two images of photometric stereo, the above described method can only be used on the parts of the surface that are illuminated in both images. Therefore the *image plane* must be defined as all points that are out of the self shadow in both images.

In order to correctly recover the height profile the 'true' normals have to be chosen.

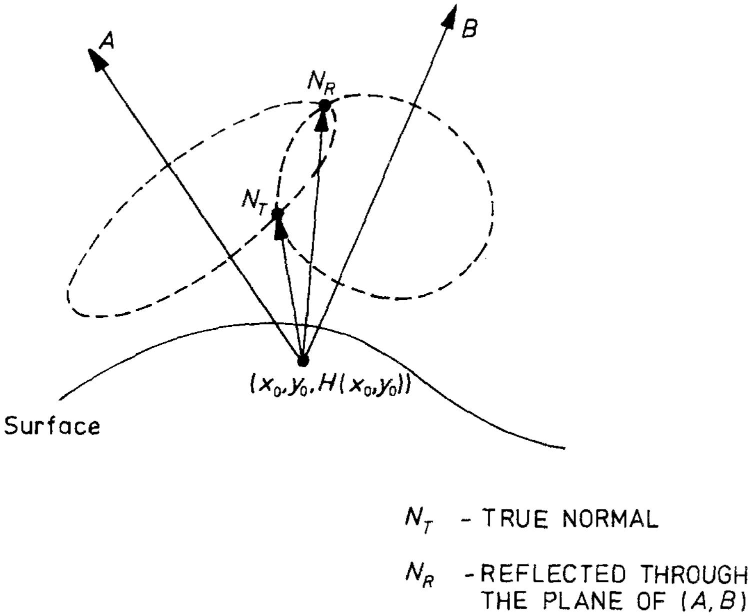

Fig. 5(a). Photometric stereo geometry.

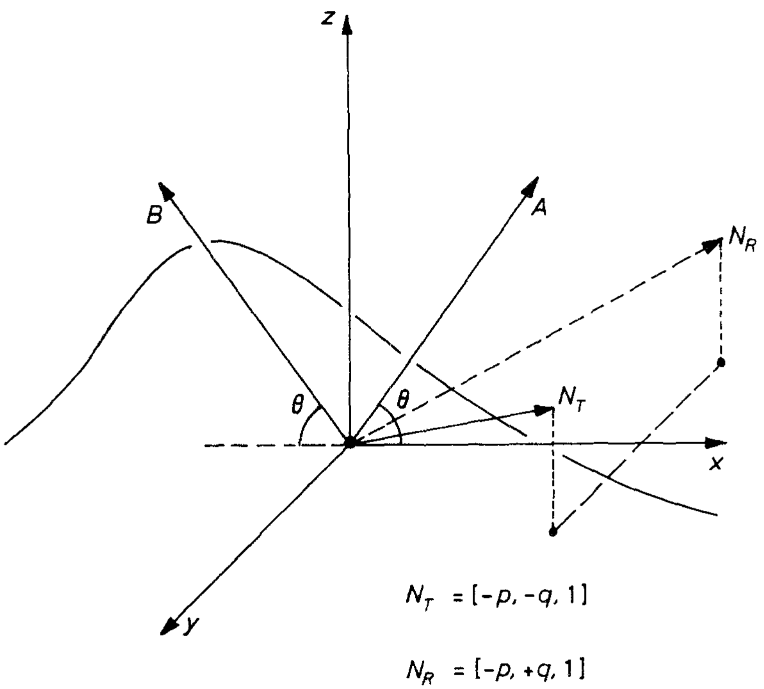

Fig. 5(b). Photometric stereo with symmetric illumination.

A trivial way to determine the true normals would be by taking yet another image under a different lighting condition, and this was indeed proposed by Woodham [W1]. However we can also exploit the lateral constraints on the normals due to the assumed continuity and smoothness of the surface.

5.2. USING THE CONTINUITY CONSTRAINT

Assuming that the normals to the height profile are continuous, consider the function $T(x, y)$, where T was defined above, i.e., $T = (1 + p^2 + q^2)^{-1/2}$. T is clearly continuous, being the continuous function of the variables p and q. Denote the two solutions of the quadratic equation (5.7), T_1 and T_2 and let the T_1 solution be defined as corresponding to the normal with a positive projection on the direction of the vector $A \times B$, T_2 being the other solution. By the discussion at the end of the previous section it is obvious that the two possible solutions will be symmetric with respect to the plane defined by the two illumination vectors A and B, unless they both collapse to a single solution situated in this plane.

Let us further define the following three sets of points in the image plane:

$$V_0 := \{\text{points where } T = T_1 = T_2\},$$

$$V_1 := \{\text{points where } T = T_1 \neq T_2\},$$

$$V_2 := \{\text{points where } T = T_2 \neq T_1\}.$$

Obviously every point on the image plane belongs to one and only one of the three sets. By continuity, the V_0 regions and the self-shadow regions divide the image plane into connected regions in each of which the normals continuously vary on the same side of the illumination-vectors-defined plane. This result follows from the following elementary

LEMMA. *Let (x_1, y_1) be a point on the image belonging to V_1, (x_2, y_2) a point belonging to V_2, and P any path from (x_1, y_1) to (x_2, y_2) where P is wholely contained in the above-defined* image plane. *Then P must contain a point belonging to V_0.*

Therefore the image plane is divided into distinct connected regions each wholely contained in one of the three sets V_0, V_1, V_2, and if we could label each region we would know the true normals everywhere on the image plane. Moreover any two regions contained in V_1 and V_2 respectively, must be separated by a region (possibly a curve) contained in V_0. The points belonging to V_0 coincide with the points where the quadratic equation (5.7) will have only one solution and its discriminant will be zero. Furthermore the assumption that the function $H(x, y)$ is smooth can be used to identify to which of the sets the points of each region belong.

5.3. USING THE INTEGRABILITY CONSTRAINT

The two functions $p(x, y)$ and $q(x, y)$ are not independent. They are connected by the fact that for a function $H(x, y)$, for which the second derivatives exist, they obey the following equation

$$\frac{\partial^2 H}{\partial x\,\partial y} = \frac{\partial^2 H}{\partial y\,\partial x}, \tag{5.8}$$

which means for p and q that

$$\frac{\partial p}{\partial y} = \frac{\partial q}{\partial x}. \tag{5.9}$$

In general, only one of the two pairs of functions (p_1, q_1) and (p_2, q_2) provided by solving (5.7) will satisfy (5.9). 'In general', here has the following meaning: (5.9) does not hold for both (p_1, q_1) and (p_2, q_2) unless the height profile satisfies some very specific constraints. These constraints are discussed below.

Suppose that a surface generates the partials p and q. Given the illumination directions A and B, and the photometric stereo data, we shall be able to determine (at each point on the surface illuminated from both directions) a pair of normals, N_t, the true normal $[-p, -q, 1]$ and a reflected normal N_r. Let us consider for simplicity that the two directions A and B are both in the plane $x - z$ being symmetric with respect to the z axis, i.e., $A = [-\sin\theta, 0, -\cos\theta]$ and $B = [\sin\theta, 0, -\cos\theta]$ for some θ. In this case the true and reflected normals will be (see Figure 5b)

$$N_t = [-p, -q, 1] \quad \text{and} \quad N_r = [-p, q, 1]. \tag{5.10}$$

This means that we must have

$$\frac{\partial}{\partial y}\, p = \frac{\partial}{\partial x}\, q = -\frac{\partial}{\partial x}\, q, \tag{5.11}$$

implying that

$$\frac{\partial^2}{\partial x\,\partial y}\, H(x, y) = 0. \tag{5.12}$$

Therefore, both choices for the surface normal, provided by the photometric stereo information, will obey the integrability condition only if the surface obeys, within some region, Equation (5.12). The general solution of this equation is easily seen to be a function of the form

$$H(x, y) = F(x) + G(y) \tag{5.13}$$

with arbitrary smooth functions $F(\cdot)$ and $G(\cdot)$. This discussion might seem to be restricted to the case of illumination directions A and B as specified above, however, we can always choose a coordinate transformation that brings us to this case, and the illuminated surface, in these new coordinates would have to satisfy (5.12), in order to have an ambiguous solutions, even when integrability is tested on both choices of normals. The coordinate transformation does not affect the shading data, which, in the Lambertian case, is independent of the position of the viewer. Therefore, we shall not be able to choose between the two normals by checking integrability, in the cases when locally, the surface can be expressed as (5.13), in the suitably defined coordinate system. An obvious example of a surface that has the form (5.13), is the case of planar surfaces. Such surfaces will remain ambiguous for all illumination directions. In general, the condition that the surface has to satisfy to remain ambiguous is seen to be very stringent, and dependent on the illumination directions. A surface of the form (5.13) will not remain such, if a coordinate transformation is performed.

Since arbitrary curved surfaces will usually not satisfy (5.13), with respect to the given directions of illumination, it can be expected that for all of the connected regions R, separated by V_0 points and/or self-shadows, only one of the following expressions will be zero

$$\int_{(x,y)\in R} \left(\frac{\partial p_1}{\partial y} - \frac{\partial q_1}{\partial x}\right)^2 \mathrm{d}x\,\mathrm{d}y, \tag{5.14a}$$

$$\int_{(x,y)\in R} \left(\frac{\partial p_2}{\partial y} - \frac{\partial q_2}{\partial x}\right)^2 \mathrm{d}x\,\mathrm{d}y \tag{5.14b}$$

Testing which of the two expressions is null, yields a labeling of the regions. If (5.14a) is zero, the pairs (p_1, q_1) are the true surface normals over region R, and the points of R belong to V_1, If (5.14b) is zero, the pairs (p_2, q_2) describe the correct surface over region R, and the points of R belong to V_2. As all the points belonging to V_0 have already been found, the pair (p, q) is determined for each point in the image plane, and we may proceed to the second part of the reconstruction. In the unfortunate but nongeneric, and hopefully rare case when some region remains ambiguous, i.e., both expressions (5.14) are zero, we shall have to check both solutions and decide which one best fits the boundary conditions provided by the neighboring regions. Height reconstruction from normals is a standard integration problem, and we shall not dwell on it here.

6. Concluding Remarks

We have presented several cases where mathematical results have found interesting applications in the new and rapidly developing field of computer vision. Some other interesting topics in vision where mathematical results have found applications are: analyzing three-dimensional object motion from their projected images, motion recovery and camera calibration from point correspondences, structured light and active vision systems and recovery of depth from line-drawing images [Kan], [Sug], [BB], [Koe] problems arising due to image digitizations over pixel-grids, [Pav], [RK], [Ser], [VG], and issues concerning qualitative rather than quantitative inference from images. Our choice of topics was, of course, heavily biased toward subjects we were actively involved in recent years, however we hope that we succeeded to convince the reader that the field of computer vision can be a rich source of problems for applications-oriented and even theoretically inclined mathematicians.

References

Local and Differential Invariants in Shape Recognition

[BHNR] Bruckstein, A. M., Holt, R. J., Netravali, A. N., and Richardson, T. J.: Invariant signatures for planar shape recognition under partial occlusion, AT&T Bell Laboratories Technical Memo, October 1991.

[BKLP] Bruckstein, A. M., Katzir, N., Lindenbaum, M., and Porat, M.: Similarity invariant recognition of partially occluded planar curves and shapes, CIS Report, Technion, June 1990, and *Artificial Intelligence and Computer Vision*, Proceedings of the Seventh Israeli AICV Conference, Elsevier, Amsterdam, 1991.

[BN] Bruckstein, A. M. and Netravali, A. N.: On differential invariants of planar curves and recognizing partially occluded planar shapes, AT&T Technical Memorandum, July 1990, and Visual Form, Capri, May 1991.

[BPB] Barrett, E. B., Payton, P., and Brill, M. H.: Contributions to the theory of projective invariants for curves in two and three dimensions, *Proc. DARPA/ESPRIT Workshop on Invariants in Computer Vision*, Iceland, March, 1991.

[COCD] Cyganski, D., Orr, J. A., Cott, T. A., and Dodson, R. J.: An affine transform invariant curvature function, *Proc. First ICCV*, London, 1987, pp. 496–500.

[IW1] Weiss, I.: Projective invariants of shapes, CAR-Technical Report, CAR-TR-339, U. of Maryland, January 1988.

[IW2] Weiss, I.: Noise resistant invariants of curves, CAR Technical Report-537, U. of Maryland, January, 1991.

[L] Lane, E. P.: *A Treatise on Projective Differential Geometry*, U. of Chicago Press, 1941.

[SB] Buchin, Su: *Affine Differential Geometry*, Gordon and Breach, 1983.

[VKO] VanGool, L., Kempenaers, P., and Oosterlink, A.: Recognition and semi-differential invariants, *Proc. 91CVPR Conference*, Hawaii, June 1991, pp. 454–460.

[W] Wilczynski, E. J.: *Projective Differential Invariants of Curves and Ruled Surfaces*, Teubner, Leipzig, 1906.

Shape Analysis via Curve Evolutions

[An] Angenent, S.: Parabolic equations for curves on surfaces, Technical Report, U. of Wisconsin, 1988.

[Ar1] Arnold, V. I.: *Geometrical Methods in the Theory of Ordinary Differential Equations*, Springer-Verlag, New York, 1988.

[Ar2] Arnold, V. I.: *Mathematical Methods of Classical Mechanics*, Springer-Verlag, New York, 1978.

[BK] Kimia, B.: Toward a computational theory of shape, PhD dissertation, Department of Electrical Engineering, McGill University, 1991.

[BSS] Bruckstein, A. M., Sapiro, G., and Shaked, D.: Affine invariant evolutions of planar polygons, CIS Report 9202, Technion, IIT, January, 1992.

[GH] Gage, M. and Hamilton, R.: The shrinking of convex plane curves by the heat equation, *J. Differential Geom.*, 23 (1986), 69–96.

[Gr] Grayson, M.: The heat equation shrinks embedded plane curves to points, I, *J. Differential Geom.*, **26** (1987), 285–314.

[Gur] Gurtin, M. E.: Thermodynamics of evolving phase boundaries, IMA Lectures Notes, University of Minnesota, 1990.

[HB] Blum, H.: Biological shape and visual science, *J. Theoret. Biol.* **38** (1973), 205–287.

[KB1] Kimmel, R. and Bruckstein, A. M.: Shape offsets via level sets, CIS Report 9204, Technion, IIT, March 1992.

[KTZ1] Kimia, B., Tannenbaum, A., and Zucker, S.: On the evolution of curves as a function of curvature, I: the classical case, to appear in *J. Math. Anal. Appl.*

[KTZ2] Kimia, B., Tannenbaum, A., and Zucker, S.: *Toward a computational Theory of shape*, Lecture Notes in Computer Science, 427, Springer-Verlag, New York, 1990.

[Lax1] Lax, P.: Shock waves and entropy, in E. Zarantonello (ed), *Contributions to Nonlinear Functional Analysis*, Academic Press, New York (1971), pp. 603–634.

[Lax2] Lax, P.: Weak solutions of nonlinear hyperbolic equations and their numerical computation, *Comm. Pure Appl. Math.* **7** (1954), 159–193.

[OS] Osher, S. and Sethian, J.: Fronts propagating with curvature dependent speed, Technical report (PAM-380), Center for Pure and Applied Mathematics, University of California, Berkeley, 1987.

[RCC] Riazanoff, S., Cervelle, B., and Chorowicz, J.: Parametrizable skeletonization of binary and multilevel images, *Pattern Recognition Lett.* **11**, (1990) 25–33.

[Se1] Sethian, J.: Curvature and the evolution of fronts, *Comm. Math. Phys.* **101** (1985), 487–499.

[Se2] Sethian, J.: An analysis of flame propagation, PhD thesis, Department of Mathematics, Univ. of Cal., Berkeley, 1982.

[SKSB] Sapiro, G., Kimmel, R., Shaked, D., and Bruckstein, A. M.: Implementing continuous scale morphology, CIS Report 9208, Technion, IIT, June 1992.

[Sm] Smoller, J.: *Shock Waves and Reaction-Diffusion Equations*, Springer-Verlag, New York, 1983.

[ST] Sapiro, G. and Tannenbaum, A.: On affine plane curve evolution, EE Publication 821, Technion, IIT, February, 1992.

Shape from Shading

[B] Bruckstein, A. M.: On shape from shading methods, *Comput. Vision, Graphics and Image Processing*, **44**, (1988), 139–154.

[BKPHB] Horn, B. K. P. and Brooks, M. J.: The variational approach to shape from shading, *Comput. Vision, Graphics and Image Processing* **33** (1986), 174–203.

[Bro] Brooks, M. J.: Surface normals from closed paths, *Proc. Sixth IJCAI*, Tokyo, Japan, 1979, pp. 98–101.

[H1] Horn, B. K. P.: Obtaining shape from shading information, in P. H. Winston, (ed), *Psychology of Computer Vision*, McGraw-Hill, New York, 1975, pp. 115–155.

[H2] Horn, B. K. P.: Hill shading and the reflectance map, *Proc. IEEE*, **69**(1), January 1981.

[H3] Horn, B. K. P.: Understanding image intensities, *Artificial Intelligence*, **8** (1977), 201–231.

[HBB] Horn, B. K. P. and Brooks, M. J. (edrs), *Shape from Shading* MIT Press, Cambridge, Mass, 1989.

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

[KB2] Kimmel, R. and Bruckstein, A. M.: Shape from shading via level sets, CIS Report 9209, Technion, IIT, June 1992.

[Max] Maxwell, J. C.: On hills and dales, *Philosophical Magazine*, London, Edinburgh and Dublin, *Journal of Science*, 1870 (pp. 233–240 in Maxwell's Collected Works).

[O] Oliensis, J.: Existence and uniqueness in shape from shading, COINS TR-89-109, Computer and Information Science, U. Mass. at Amherst, October 1989.

[Sax] Saxberg, B. V. H.: A modern differential geometric approach to shape from shading, PhD Thesis, AI Lab, MIT, 1989.

Photometric Stereo

[BKPH] Horn, B. K. P.: *Robot Vision*, MIT Press, Cambridge, Mass, 1986.

[I] Ikeuchi, K.: Determining a depth map using dual photometric stereo, *Internat. J. Robotics Res.* **6**, Spring 1987.

[K] Kozera, R.: Existence and uniqueness in photometric stereo, PhD Dissertation, The Flinders University of South Australia, Adelaide, June 1991.

[LB] Lee, S. and Brady, M.: Integrating stereo and photometric stereo to monitor the development of glaucoma, *Image and Vision Computing* **9**, (February, 1991), 39–44.

[M] Marr, D.: *Vision*, W. H. Freeman, San Francisco, 1982.

[OB] Onn, R. and Bruckstein, A. M.: Integrability disambiguates surface recovery in two-image photometric stereo, *Internat. J. Comput. Vision* **5**(1) (1990), 105–113, see also On photometric stereo, Technion EE Report 639, July 1987.

[RJW] Woodham, R. J.: Photometric Method for Determining surface orientation from multiple images, *Optical Engineering* 19 (1980).

Additional General References

[BB] Ballard, D. H., Brown, C.M.: *Computer Vision*, Prentice-Hall, Englewood Cliffs, 1982.

[Kan] Kanatani, K.: *Group Theoretical Methods in Image Analysis*, Springer-Verlag, New York, 1990.

[Koen] Koenderink, J.: *Solid Shape*, MIT Press, Cambridge, Mass., 1990.

[Pav] Pavlidis, T.: *Algorithms for Graphics and Image Processing*, Computer Science Press, 1982.
[RK] Rosenfeld, A. and Kak, A.: *Digital Image Processing*, 2nd Edn, Academic Press, New York, 1982.
[Ser] Serra, J.: *Image Analysis and Mathematical Morphology*, Academic Press, New York, 1982.
[Sug] Sugihara, K.: *Machine Interpretation of Line Drawings*, MIT Press, Cambridge, Mass., 1986.
[VG] *Vision Geometry, Contemporary Mathematics* Volume 119 (Melter, Rosenfeld and Bhattacharya, Editors), American Mathematical Society, 1991.