

SURFACE-NORMALS FROM CLOSED PATHS

M.J. Brooks
Department of Computer Science,
University of Essex,
Colchester, England.

It is now well established that the shading present in an image of an object is a potential source of 3-D shape information. With this in mind, we seek constraints arising out of the assumption that an image portrays an object with a smooth surface. On obtaining a threshold and a depth-invariance constraint (the latter requiring integration around closed curves) we offer a method of shape from shading for the purpose of demonstrating their potential usefulness.

KEY WORDS

Shape from shading, surface-normal, closed path, constraints, depth-invariance.

TOPICS

Vision and image understanding systems.

1. INTRODUCTION

Given an image of a smooth surface, we consider the problem of assigning to each image point an orientation corresponding to the appropriate object surface-normal, thus obtaining a surface-normal map. This constitutes the core of the shape from shading problem (see [3] and [8]) since 3-D surface data can be recovered from such a map by integration.

Two basic stages can be isolated in the generation of a surface-normal ($s - n$) map:

(1) The luminance value at each image point defines an orbit in the reflectance map [4]. Thus, for each image point, we produce an infinite, one-dimensional set of surface-normals.

(2) From each of these sets, we select a particular surface-normal such that the resulting surface-normal map represents a smooth surface.

Given a knowledge of the illumination conditions and surface reflectivity, the first stage is relatively straightforward. It is the second stage on which we shall concentrate.

This work is funded by the Science Research Council.

We contend that the assumption that an image depicts a smooth surface is one that has not previously been fully exploited. (For instance, Woodham's program [8] requires substantial clues as to the object shape before being able to run.) Thus we seek constraints imposed by this assumption on dense $s-n$ maps: that is, maps which have an infinitely fine resolution (for we can then use continuous mathematics). These constraints are subsequently translated to discrete form prior to application.

2. CONSTRAINING A DENSE SURFACE-NORMAL MAP

2.1 Smoothness

A dense $s-n$ map, u , represents a smooth surface if, and only if, for any point s in the map and $\delta > 0$, $\exists \epsilon > 0$ such that, for any s' in an ϵ -neighbourhood of s , the angle in radians between $u(s)$ and $u(s')$ is less than δ . (See Fig. 1.)




Fig. 1 Smoothness definition

2.2 Depth-invariant Loops

Suppose we were to draw a closed loop, C1, on a continuous surface; we would clearly finish the loop at the same point in space at which we started. Were we now to project the surface orthogonally to form a dense s-n map, then providing all points on the surface were visible with respect to the projection, we would obtain a corresponding closed loop, C2, in the resulting s-n map. If we integrated around C2 over surface-normals to recover depth, then on completing a tour of C2, we would result in no depth change. Note the similarity between this and the scene analysis concepts of spatial closure expounded in [5] and [2].

This observation leads to a formal constraint: for a dense s-n map, $u(s)$, to represent a continuous surface, the relation

$$\oint_{\Gamma} \frac{u(s) \cdot ds}{u(s) \cdot D} = 0$$

must hold for all closed paths Γ in the map. Here, s is a point on Γ with unit derivative ds as shown in Fig. 2; D is a unit vector of depth along the direction of projection. For the derivation, see [1]. The above definition has an analogue in the complex analysis theorem of Morera's; see, for example, p.538 of [6].




Fig. 2 Integrating around a curve Γ

3. MAKING THE CONSTRAINTS DISCRETE

Computer vision programs operate on tessellated images and so we require the discrete form of the constraints obtained in section 2.

3.1 Thresholding

In attempting to import the smoothness definition of 2.1 into the discrete world, we are, in a

strict sense, doomed to failure since any discrete s-n map can represent a smooth surface providing the surface is allowed to fluctuate sufficiently between the sampling points. Due to a lack of information, we cannot be expected to recover this surface. Indeed, if we are expected to recover the object shape depicted in an image, we are justified in insisting that the image capture the characteristics of the surface. For our purpose, that criterion will be fulfilled if a piece-wise planar solution (one surface-normal per image point) is a satisfactory approximation to the surface.

Suppose we were to generate several images of varying resolutions from a wildly fluctuating surface. If we produced a discrete s-n map for each image and computed, for each map, the maximum angle between any two neighbouring surface-normals, we would find that, in general, this angle decreased as the resolution increased. So by assuming a minimum level of resolution per object scale and per maximum surface curvature, we place an upper limit, or threshold, on the permissible angular difference between neighbouring normals. (The author has recently improved considerably on a blanket threshold by adopting a functional threshold dependent on surface slant.)

3.2 The two by two Integral

We now apply the depth-invariant integral to the discrete case of a 2*2 window of orientations. For a square path of integration (a 2*2 path) and surface-normal quartet as shown in Fig. 3, then for this window to be representative of a continuous surface, the relation

$$\frac{(a_1 + b_1)}{c_1} + \frac{(a_2 - b_2)}{c_2} - \frac{(a_3 + b_3)}{c_3} - \frac{(a_4 - b_4)}{c_4} = 0^*$$

must hold [1].




Fig. 3 The two by two integral

* The author has recently become aware of [9] in which an equivalent formula is (independently) derived.

4. A SHAPE FROM SHADING METHOD

The need to investigate the depth-invariance of all closed paths in a discrete s-n map is an intolerable burden which we finesse by using the following result.

Lemma

If all 2×2 paths are depth-invariant in a discrete s-n map, then so too is any closed path in the map.

Consider a dense s-n map that represents a continuous surface. Then the depth-integral around a closed curve Γ clearly equals the sum of the depth-integrals around the constituent closed curves Γ_1 and Γ_2 as shown in Fig. 4.




Fig. 4 Constituent curves

This splitting can be repeated as often as required and providing the correct directionality is observed, the sum of the depth-integrals around the constituent curves will equal the depth-integral around Γ . For a complex analysis analogue, see Goursat's proof, on p.526 of [4].

Now, any closed path in a discrete s-n map will consist of linked horizontal and vertical sections (e.g., Fig. 5). But any such path can be dissected into constituent paths that are all 2×2 (e.g., Fig. 6). Further, if these 2×2 paths all integrate to zero in either direction, then so too, by the above, must the bounding closed path.

We now describe a method of shape from shading in order to demonstrate the potential of these constraints. The method is in two stages.




Fig. 5 A discrete s-n map closed curve




Fig.6 Constituent 2×2 paths

4.1 Obtaining a Set of Interpretations for each 2×2 window

As mentioned earlier, each image point specifies, in general, an infinite, one-dimensional subset of all possible surface-normals. Crucial to our method is the quantising of this interval into a finite set of orientations.

We therefore have a potentially vast number of possible combinations of surface-normal interpretations for each 2×2 window, of which the majority are likely to fail either the threshold or depth-invariance test. There is insufficient space here to detail a speedy means of generating a set of legal interpretations, but note that there is considerable scope for Waltz-like filtration on neighbouring surface-normals (see [7]). Also, note that all 2×2 windows can be processed in parallel. This stage has recently been successfully implemented.

4.2 Grow and Encompass

Having generated a set of legal surface-normal interpretations for each 2×2 window, we now want to obtain a set of interpretations for the whole image. To do this, we start by taking a 2×2 window of interpretations and merging this with an intersecting 2×2 window to form a larger area of legal interpretations. We continue to expand this area until the whole image is covered, thereby obtaining a set of solution s-n maps.

We define the algorithm inductively by considering the merging of the interpretations of an n-pixel area and a partially intersecting 2×2 area.

Suppose, in the n-area, pixels are indexed from 1 to n, and that the 2×2 area intersects the n-area at, say, two image points. Let $(p_{m1}, \dots, p_{mn})$ be a legal interpretation of the n-area such that p_{mi} is the orientation of the i^{th} pixel, $1 \leq i \leq n$. Let $(p_{m1}, \dots, p_{mn})$ for $m = 1, \dots, k$ be the set of legal interpretations for the n-area, denoted as $(p_1, \dots, p_n)_k$. Let the legal interpretation set for the 2×2 window be $(q_a, q_b, q_{n+1}, q_{n+2})_l$ where $a \neq b$ and $a, b \in [1, n]$.

Now, for all

$$(p_{m1}, \dots, p_{mn}) \in (p_1, \dots, p_n)_k$$

and

$$(q_a, q_b, q_{n+1}, q_{n+2}) \in$$

$$(q_a, q_b, q_{n+1}, q_{n+2})_l$$

we test the relation

$$(q_{ra} = p_{ma} \text{ AND } q_{rb} = p_{mb})$$

which checks whether the surface-normals at the points of intersection are equal. If this relation is false, we pass on to the next combination, thereby rejecting the merger, otherwise we output

$$(p_{m1}, \dots, p_{mn}, q_{rn+1}, q_{rn+2})$$

as a legal interpretation of the resultant area.

5. CONCLUSION

We have outlined a shape from shading algorithm that converges on legal interpretations without needing clues as to the solution shapes. This was done merely to demonstrate the value of the constraints in section 3. The labourious behaviour of the algorithm is clearly inadequate and the author is currently implementing a

cooperative algorithm.

ACKNOWLEDGEMENTS

I am grateful to Pat Hayes, Mike Brady, Tony Cohn, Steven Draper and Paul Robertson for their useful comments and to Berthold Horn for his valuable encouragement.

REFERENCES

- [1] Brooks, M.J. "Miscellaneous results", unpublished, 1979.
- [2] Draper, S. "The penrose triangle and a family of related figures", Perception, 1978, vol. 7, p.283-296.
- [3] Horn, B.K.P. "Obtaining shape from shading information", in The Psychology of Computer Vision, P.H. Winston, ed., McGraw-Hill, New York, 1975, p.115-155.
- [4] Horn, B.K.P. "Understanding image intensity", Artif. Intell. 8 (April 1977), p.201-231.
- [5] Huffman, D.A. "A duality concept for the analysis of polyhedral scenes", Machine Intelligence 8, eds. B. Meltzer, D. Michie, Edinburgh University Press, p.475-492.
- [6] Kreyszig, E. Advanced Engineering Mathematics, Wiley, 1972.
- [7] Waltz, D. "Understanding line drawings of scenes with shadows", in The Psychology of Computer Vision, P.H. Winston, ed., McGraw-Hill, New York, 1975.
- [8] Woodham, R. "A cooperative algorithm for determining surface orientation from a single view", in IJCAI 1977, p.635.
- [9] Strat, T.M. "A numerical method for shape-from-shading from a Single Image", Ph.D. thesis, MIT, 1979.