Curved Mondrians: shading analysis of patterned objects[1]

WALTER F. BISCHOF

Department of Psychology and Alberta Centre for Machine Intelligence and Robotics, University of Alberta, Edmonton, Alta., Canada T6G 2E9

AND

MARIO FERRARO

Istituto di Fisica Superiore, Università di Torino, Torino, Italy

Received September 20, 1988

Revision accepted September 25, 1989

Most shape-from-shading methods assume that surface reflectance is constant within large image regions. This assumption is violated in natural scenes with objects made from different materials. We present a more general method for recovering shape from shading, assuming that surfaces are smooth and albedo is piecewise constant, as would be the case if a Mondrian image was painted on a smooth curved surface. Our method is based on combining Brooks and Horn's method for shape recovery with the recovery of albedo using stochastic relaxation.

Key words: shape from shading, stochastic relaxation, Mondrian image.

La plupart des méthodes de la figure dérivée de l'ombre tiennent pour acquis que le pouvoir réflecteur d'une surface est constant à l'intérieur d'importantes surfaces d'image. Cette hypothèse est contredite dans les scènes naturelles avec des objets fabriqués de matériaux différents. Nous présentons une méthode plus générale de récupération de la figure à partir de l'ombre, qui suppose que les surfaces sont lisses et que l'albédo est constant, comme ce serait le cas si une image de Mondrian était peinte sur une surface courbe lisse. Cette méthode combine la méthode de récupération de la figure de Brooks et Horn et celle de la récupération de l'albédo à l'aide de la relaxation stochastique.

Mots clés : figure dérivée de l'ombre, relaxation stochastique, image de Mondrian.

[Traduit par la revue]

Comput. Intell. 5, 121–126 (1989)

1. Introduction

Computational vision aims at understanding how three-dimensional representations of the world can be reconstructed from information contained in two-dimensional images. Recent research has produced a variety of methods that allow the recovery of such shape information, including, for example, shape-from-stereo, shape-from motion, shape-from-texture, and shape-from-shading. Most of these shape recovery problems are ill posed (Hadamard 1923; Torre and Poggio 1986) in that existence, uniqueness, and stability of solutions may not be guaranteed in the absence of additional constraints. In the case of shape-from-shading, for example, Pentland (1984) restricts the space of solution surfaces to umbilical surfaces to enable unique recovery of surface orientation from local variation of image irradiance. Brooks and Horn (1985) use a weaker constraint, requiring that the solution surface maximizes some global smoothness measure, but constraints on surface shape have to be propagated nonlocally. Both approaches are based on the image irradiance equation

$$[1] \quad E(x, y) = \rho\lambda n(x, y) \cdot s$$

where E is the image irradiance, ρ the surface albedo, λ the incident flux, n the surface normal, and s the illuminant direction (sun). They require that surface reflectance, in the imaging model [1] Lambertian reflectance, remains constant over a large area, and all variations in image irradiance are attributed to variations in surface orientation. Further, the two approaches provide no means for detecting that this assumption may be violated, except in some pathological cases. Such is the case, for example, with a black area in an image with $s = v$, the viewing direction, where surface orientation is computed being orthogonal to the viewing direction over some extended area, a geometrically impossible inference.

The aim of this paper is to investigate the feasibility of shape inference under weaker constraints than those imposed by either Pentland or Brooks and Horn. There is no way of separating the two effects of surface orientation and surface reflectance in the general case, since any change in image irradiance can be attributed to a change in either of the two or a combination of both. Thus a photograph can be interpreted as a flat surface with changing reflectance or as an image of a curved surface with or without changes in reflectance. However, we will show that surface orientation can be successfully recovered if the surface is assumed to be smooth (as in Brooks and Horn) and surface albedo is piecewise constant (as opposed to globally constant in Brooks and Horn).

Piecewise constant surface albedo can lead to discontinuities in image irradiance and thus the borders of "albedo patches" could be detected by either looking for discontinuities in image irradiance directly (edge detection) or looking for large errors in the predicted image irradiance after a surface orientation fit, assuming constant albedo. As we will show later, both these approaches lead to unsatisfactory results, inferior to our scheme in which surface orientation and surface albedo are recovered simultaneously.

Our approach is closely related, on one hand, to Brooks and Horn (1985) for recovering surface orientation and, on

[1]A previous version of this paper appeared in the Proceedings of the Canadian Society for the Computational Studies of Intelligence (CSCSI-88) Conference, Edmonton, Alberta.

the other hand, to Marroquin *et al.* (1987) and Geman and Geman (1984) for recovering the piecewise constant albedo map. As discussed in the next section, the two processes run in parallel, each updating iteratively the input to the other process.

2. Model

We want to recover the orientation, $n(x, y)$, of a smooth surface satisfying the image irradiance equation for Lambertian surfaces

$$[2] \quad E(x, y) = \rho(x, y)n(x, y) \cdot s$$

over some domain Ω, where $E(x, y)$ is the image irradiance, $\rho(x, y)$ the (piecewise constant) surface albedo, $n(x, y)$ the unit normal of the surface, and s the (known) direction of a distant single point source (sun) with $|s| = \lambda$, the (known) incident flux. It is obvious that we put quite strong restrictions into the image irradiance equation, such as the restriction to Lambertian surfaces. We have done so in order to allow a direct comparison with experiments reported in Brooks and Horn (1985). It is clear, however, that these restrictions have to be relaxed if our method is to be applied to natural images and that it has to be complemented, for example, with methods for dealing with specularities (Healey and Binford 1986; Brelstaff and Blake 1988).

In the remainder of this section we will first discuss how to recover the surface orientation map, $n(x, y)$ (Sect. 2.1), then how to recover the surface albedo, $\rho(x, y)$ (Sect. 2.2), and then how to combine the two methods (Sect. 2.3). It should, however, be pointed out from the beginning that in our implementation the two processes for recovering orientation annd albedo are not operating independently or sequentially, but are intertwined in such a way that in each iteration the output of one process is used as input for the other process.

2.1. Recovering the surface orientation map

As shown in Brooks and Horn (1985), recovering the surface orientation map, $n(x, y)$, from [2] can be put as a variational problem where we try to minimize [3] with respect to n.

$$[3] \quad I(n) = \iint_\Omega \{[E(x, y) - \rho(x, y)\, n(x, y) \cdot s]^2 + \alpha[n_x^2(x, y) + n_y^2(x, y)] + \mu(n^2 - 1)\}\, dxdy$$

In [3] the first term in the integral captures errors of the predicted image irradiance, the second term captures deviations from smoothness of surface orientation, with α weighting the relative importance of this term and n_x and n_y denoting partial derivatives, and the third term constrains $n(x, y)$ to unit vectors. In terms of regularization theory (Tikhonov and Arsenin 1977; Torre and Poggio 1986), the second term is called the stabilizing functional and α the regularization parameter.

The deficiencies of the smoothing term in [3] are well known (cf. Horn and Brooks 1986; Frankot and Chellappa 1987), but as was pointed out earlier, adopting this form allowed a comparison with the results reported in Brooks and Horn (1985). Disadvantages of this form of the smoothing term are further discussed below. The minimization formulation [3] differs from that of Brooks and Horn in that the term for surface albedo, $\rho(x, y)$, is not assumed to be (globally) constant. As will be discussed below, iterative estimates for $\rho(x, y)$ are provided by the process that recovers the surface albedo map. The Euler–Lagrange equation associated with [3] is

$$[4] \quad (E - \rho n \cdot s)\rho s + \alpha \nabla^2 n - \mu n = 0$$

where

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

is the Laplacian. Using the 9-point discrete approximation for the Laplacian

$$[5] \quad \nabla^2 \approx \frac{1}{6\epsilon^2} \begin{bmatrix} 1 & 4 & 1 \\ 4 & -20 & 4 \\ 1 & 4 & 1 \end{bmatrix}$$

and solving [4] for n, we arrive at the following iterative estimates for n:

$$[6] \quad n_{ij}^{k+1} = m_{ij}^{k+1}/|m_{ij}^{k+1}|$$

$$[7] \quad m_{ij}^{k+1} = \bar{n}_{ij}^{k} + \frac{3\epsilon^2}{10\alpha}(E_{ij} - \rho_{ij} n_{ij}^{k} \cdot s)\rho_{ij} s$$

where m_{ij}^{k+1} are the unnormalized surface normals and

$$[8] \quad \bar{n}_{ij}^{k} = \frac{1}{20}[4(n_{i,j+1} + n_{i,j-1} + n_{i+1,j} + n_{i-1,j}) + (n_{i+1,j+1} + n_{i+1,j-1} + n_{i-1,j+1} + n_{i-1,j-1})]$$

After every iteration of the surface orientation estimation, we can obtain an estimate of the surface albedo, $\hat{\rho}(x, y)$, from [2]:

$$[9] \quad \hat{\rho}(x, y) = E(x, y)\, (n(x, y) \cdot s)^{-1}$$

except at self-shadow boundaries where $n(x, y) \cdot s = 0$.

In the experiments reported in Sect. 3, the surface orientation, n, was initialized to the correct value at the occluding boundary and to $n = [0, 0, 1]$ elsewhere. Albedo $\rho(x, y)$ was initialized to a random value in the first iteration.

2.2. Recovering the surface albedo map

Given our assumption that albedo be piecewise constant, we want to fit a piecewise constant albedo map, $\rho(x, y)$, to the estimates $\hat{\rho}(x, y)$ obtained in [9]. That is, $\rho(x, y)$ should be constant except at boundaries between different regions, and the discontinuity boundaries should be spatially continuous. This problem could again be cast in terms of a minimization problem, but it is unlikely that the associated minimization functional is convex and thus could be solved using standard variational methods.

One way to overcome the problem of finding global minima of nonconvex functions is to use simulated annealing (Kirkpatrick *et al.* 1983), an extension of the algorithm of Metropolis *et al.* (1953) for simulating statistical mechanical systems. In this method, randomness is introduced into the steepest-descent path in order to avoid being trapped in local minima. The degree of randomness is controlled by a temperature parameter, which is initially high and allowed to cool down as the system approaches the global minimum.

Simulated annealing and its relatives have been successfully applied to a variety of problems including the restoration of noise-corrupted piecewise constant images (Geman and Geman 1984; Marroquin *et al.* 1987). Our approach to recovering the piecewise constant albedo map, $\rho(x, y)$, from

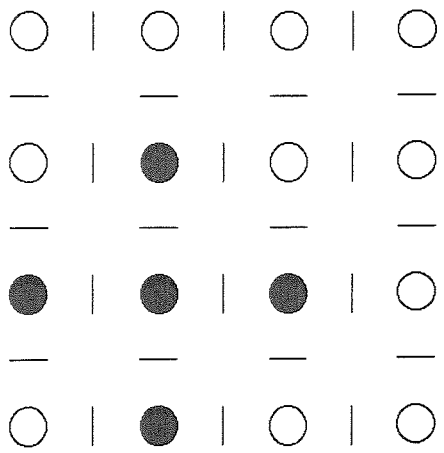

FIG. 1. Albedo node lattice with interstitial (horizontal and vertical) line lattice for modelling discontinuity boundaries. The black nodes indicate an "albedo node" with its four neighbour nodes.

the albedo estimates, $\hat{\rho}(x, y)$, follows closely their approach. The goal of the recovery process is to determine the "true" albedo value at every node of a dense grid in the viewer coordinate system. As a side effect, the method also produces the albedo discontinuity boundaries which may be located between any pair of neighbouring albedo nodes in a 4-neighbourhood system (see Fig. 1).

In the discussion of the albedo map recovery, we first introduce the associated energy functions and then discuss the recovery process using simulated annealing.

2.2.1. Energy functions

To model the piecewise constant albedo map, we use a coupled node-line model (Geman and Geman 1984; Marroquin *et al.* 1987) with the node process capturing penalties on albedo values and the line process capturing penalties on the local geometry of discontinuity boundaries.

• *Node process.* At every node in the lattice the modelled albedo, ρ_i, should be as close as possible to the estimated albedo, $\hat{\rho}_i$, obtained from [9], and the albedo value between neighbouring nodes should not differ unless there is a discontinuity boundary between them. To model the former we introduce the "albedo error" energy, D_ρ, and to model the latter we introduce the "albedo variation" energy, V_ρ. The *albedo error energy* for node i is defined as

$$[10] \quad D_\rho(i) = |\rho_i - \hat{\rho}_i|$$

We used the absolute error as opposed to a more usual quadratic error term in [10] based upon experiments with our method. Using a quadratic term, the albedo error tended too often to dominate the combined energy function introduced below.

The *albedo variation energy* is defined as follows. Let i and j be two neighbouring nodes, l_{ij} the line element between the two nodes, and N_i the set at all neighbours (the clique) of node i. Then the albedo variation energy, $V_\rho(i)$, is defined as

$$[11] \quad V_\rho(i) = \sum_{j \epsilon N_i} V'_\rho(i, j, l_{ij})$$

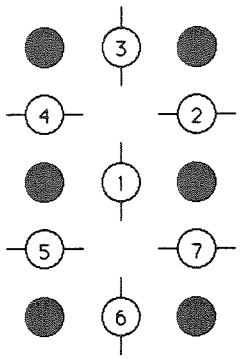

FIG. 2. Vertical line element (1) with six neighbours (2–7). Energy functions are defined on the two cliques, (1, 2, 3, 4) and (1, 5, 6, 7).

where

$$[12] \quad V'_\rho(i, j, l_{ij}) = \begin{cases} -1 & \rho_i = \rho_j,\ l_{ij} = \text{"off"} \\ +1 & \rho_i \neq \rho_j,\ l_{ij} = \text{"off"} \\ 0 & l_{ij} = \text{"on"} \end{cases}$$

where l_{ij} = "on" or "off" indicates the presence or absence of a discontinuity between nodes i and j. The assumption that the surface albedo is piecewise constant is captured in the term V'_ρ. For any two neighbouring nodes, equal albedo values within regions are favoured over equal albedo values across boundaries. Unequal albedo values are slightly penalized across region boundaries (l_{ij} = "on"), with the effect that globally constant albedo is preferred over piecewise constant albedo. Finally, unequal albedo values within a region (l_{ij} = "off") are strongly penalized. Note that under the assumption of piecewise constancy any difference in albedo, independent of the amount, should be equally penalized. Finally, note that there is no coupling between [12] and the data term [10], since the albedo variation term, [12], is defined between nodes whereas the albedo error term, [10], is defined for a single node only. (In this respect, [12] differs from formulations found, for example, in Poggio *et al.* (1988) or Hutchinson *et al.* (1988).)

• *Line process.* In modelling the spatial geometry of discontinuity boundaries, isolated discontinuities and clustered discontinuities should be highly penalized, whereas "good continuations" of boundaries should be favoured. Figure 2 shows a vertical line element (1) with six neighbouring line elements (2–7), defining the neighbourhood of element (1), and forming two cliques, (1, 2, 3, 4) and 1, 5, 6, 7). The energy functions, $V_l(l_{ij})$, are defined for each clique with the values of $V_l(l_{ij})$ for all possible configurations, up to rotations, given in Fig. 3. Isolated boundary fragments (V_2) and clusters of boundary fragments (V_5 and V_6) are highly penalized.

Given some estimated albedo values, $\hat{\rho}_i$, we want to find the interpretation, $I = (P, L)$, of albedo values $P = \{\rho_i\}$ and discontinuity boundaries, $L = \{l_{ij}\}$, which minimizes the combined energy function, $U(I, L)$:

$$[13] \quad U(I, L) = \beta_1 \sum_{i \epsilon I} D_\rho(i) + \beta_2 \sum_{i \epsilon I} V_\rho(i) + \beta_3 \sum_{l \epsilon L} V_l(l)$$

where β_1, β_2, and β_3 are weighting factors.

In the experiments reported in Sect. 3, the initial conditions were set as follows: Albedo values $\rho(x, y)$ were initialized to a random value. Discontinuity line elements were

initialized randomly (with probability $p = 0.5$) to "on" or "off", except at the occluding boundary where they were all set to "on".

2.2.2. Simulated annealing

Given the albedo estimates, $\hat{\rho}(x, y)$, we want to find the values of $\rho(x, y)$ minimizing the combined energy function [13]. As was pointed out before, there is no guarantee that the function [13] is convex and thus steepest-descent algorithms may be trapped in local minima. There exist several methods to overcome this problem, such as the graduated nonconvexity algorithm (Blake and Zisserman 1987) or a dynamic programming algorithm (Papoulias 1985). The method we use for the minimization of [13] is based on the Metropolis algorithm (Metropolis *et al.* 1953) and on the annealing algorithm (Kirkpatrick *et al.* 1983). It can be briefly described as follows. Instead of using a deterministic scheme for updating albedo values, the following scheme is used. New values of $\rho(x, y)$ are chosen randomly and the difference in energy, $\Delta U = U_{\text{new}} - U_{\text{old}}$, is computed. A new value of $\rho(x, y)$ is accepted if $\Delta U \leq 0$, otherwise it is accepted with probability $p = \exp(\Delta U/T)$ with the parameter T being called temperature. When T is large, energy increasing changes are often accepted and local minima can be escaped. As $T \to 0$ the system "freezes" and an almost deterministic updating scheme is followed. The temperature, T, is lowered according to the schedule $T = T_0/\text{ld}(1 + k/4)$ with k being the iteration number and ld being the logarithm to the base 2. Geman and Geman (1984) show that such a schedule for lowering temperature is suf ficient for converging to states at least close to the global minimum.

After updating all albedo values, ρ, the same process is repeated for the introduction or removal of line elements, l (see Geman and Geman 1984). It should be noted that in the original annealing method (Kirkpatrick *et al.* 1983), temperature is not changed until the average energy reaches equilibrium. We change temperature at every iteration regardless of whether equilibrium has been reached, but compensate for it by using a conservative cooling scheme.

2.3. Combining the recovery of surface orientation and surface albedo

There are several approaches to combining the recovery of surface orientation and albedo. In this section we first present the approach we used in our implementation, simultaneous recovery, and then discuss two alternative approaches involving sequential recovery.

The approach we have found most successful involves the simultaneous recovery of surface orientation and surface albedo. Surface orientation recovery using the Brooks and Horn scheme is intertwined with albedo recovery using simulated annealing. In every iteration the surface orientation process produces an estimate of surface albedo, $\hat{\rho}(x, y)$, which is used by the albedo process to produce an estimate of the piecewise constant albedo map, $\rho(x, y)$. The new albedo map is then used in the next iteration of the surface orientation process. Simultaneous recovery of surface orientation and surface albedo has clear advantages over approaches that attempt to recover the two sets of unknown sequentially.

Consider first the "albedo-surface" approach in which albedo, or at least boundaries of patches with constant albedo, are recovered first. Given the model assumption of surface smoothness, discontinuities in image irradiance can only occur in two ways, at points with a discontinuity in albedo and at occluding boundaries of surfaces. The latter are assumed to be known in the model for the recovery of surface orientation. Hence boundaries of (constant) albedo regions could be detected using some edge detection mechanism. If this can be done reliably, then a mechanism for orientation recovery could be devised, which independently estimates surface albedo in each of the regions. Such a scheme faces several difficulties. First, the reliability of edge detection decreases in regions near self-shadow boundaries where image irradiance changes rapidly, violating our assumption that sharp changes in image irradiance must be due to albedo changes. Second, the subsequent estimation of albedo for each region requires that the region boundaries obtained by the edge detection mechanism are closed, a requirement that cannot be easily guaranteed. Accordingly, our experiments using only an edge detection scheme proved not sufficient for recovering boundaries of albedo regions reliably. However, information from the edge detection mechanism could be used, for example, for initializing line elements in the albedo recovery mechanism.

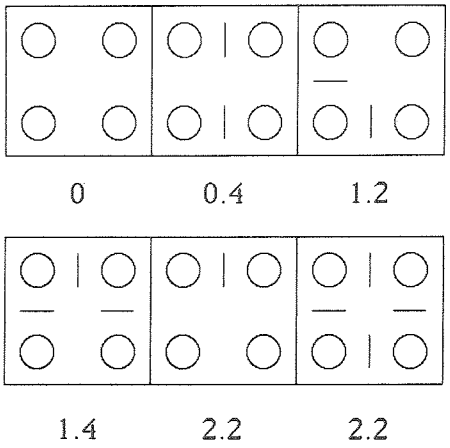

FIG. 3. Six possible line configurations and their associated energies.

Consider next the "surface-albedo" approach in which boundaries of albedo patches are detected after a full cycle of surface orientation recovery with a "constant albedo" model. Albedo boundaries can then be located using some discontinuity detector (Terzopoulos 1985), for example, by detecting positions where the first term in [3] has a significantly large value. Given these boundaries and some appropriate albedo estimation for each region, surface orientation can then be re-estimated. One problem with this approach is that the reliability of the discontinuity detector is reduced by the previous smooth surface interpolation.

3. Experiments

Recovery of surface orientation and albedo was tested on synthetic, noise-free images of size 64^2 and with 2^8 intensity levels. Illuminant direction, s, was coincident with viewing direction, v, and incident flux was kept constant, $\lambda = |s| = 1$. Images and results of the recovery process are shown in Fig. 4. The first surface (Fig. 4*a*) was a Mondrian-like sphere with seven patches of different albedo in the range 0.2–1.0, the second surface (Fig. 4*b*) was an egg-shaped Mondrian surface, and the third surface (Fig. 4*c*)

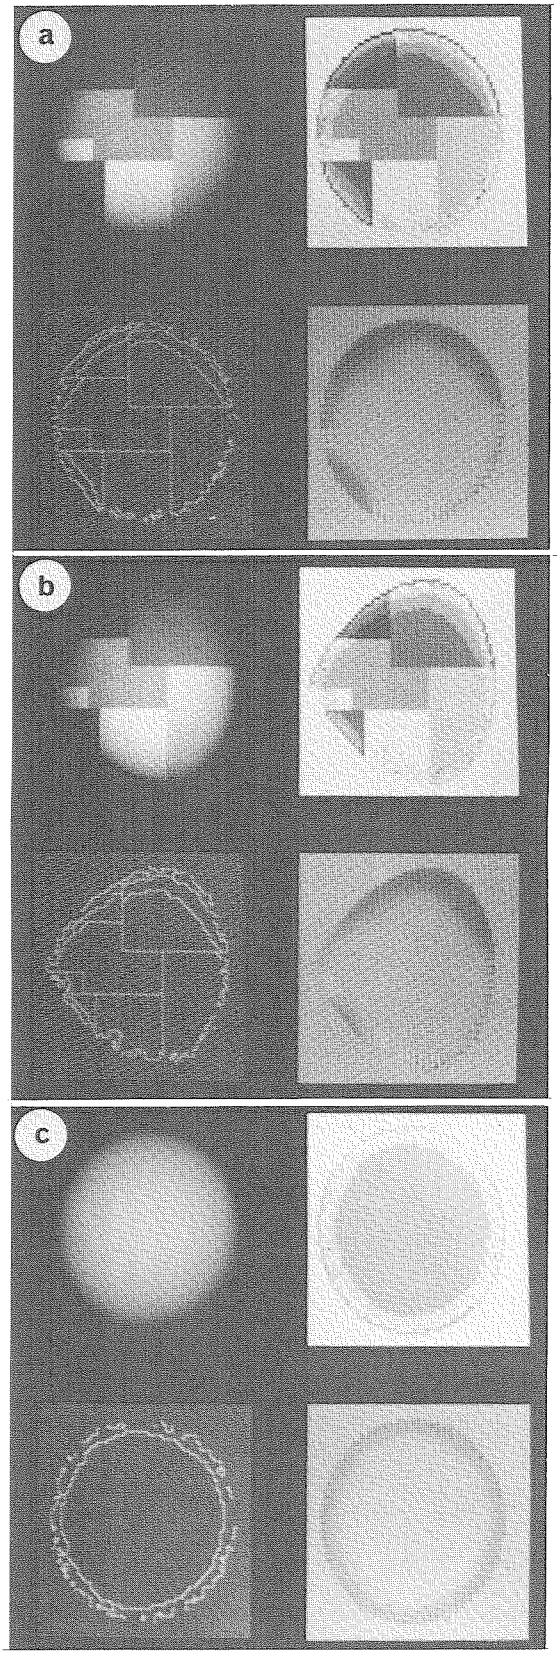

FIG. 4. Images and results for albedo and shape recovery for (*a*) a sphere with piecewise constant albedo, (*b*) an egg-shaped object with piecewise constant albedo, and (*c*) a sphere with constant albedo. The images are shown in the upper left, the recovered albedo map in the upper right, and the inferred discontinuity boundaries in the lower left. The error of the *z*-component ($n \cdot s$) of the recovered surface orientation is shown in the lower right. Grey areas indicate correct recovery, in dark areas recovered surface is too steep, and in white areas too flat.

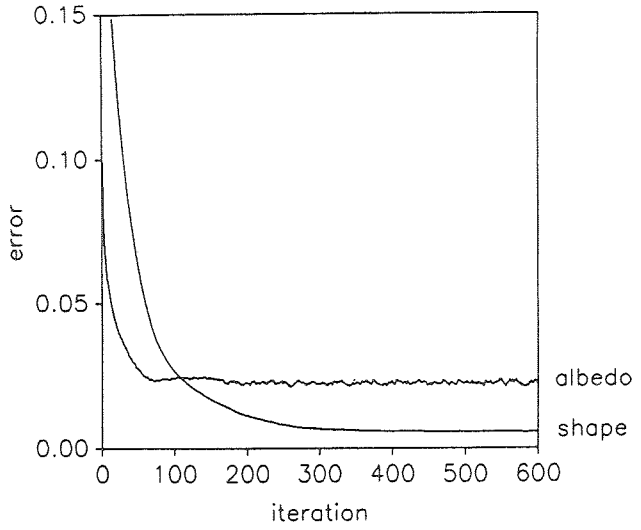

FIG. 5. Average error of estimated surface orientation, $E[n_{\text{inferred}} \cdot n_{\text{true}}]$, and average error of estimated albedo, $E[\rho_{\text{inferred}} - \rho_{\text{true}}]$, over 600 iterations for egg-shaped surface with piecewise constant albedo.

— used for control purposes — was a sphere with constant albedo.

The images are shown in the upper left, the inferred albedo map in the upper right, the inferred discontinuity boundaries in the lower left, and the error of the inferred surface orientation in the lower right of Figs. 4*a*–4*c*. The latter is to be interpreted as follows: in grey areas the *z*-components ($n \cdot s$) of the inferred surface normals are correct, in white areas too high (inferred surface orientation too flat), and in black areas too low (inferred surface orientation too steep).

All examples were computed using the same parameter values: $\alpha = 0.3$ (see [3]), $\beta_1 = 0.05$, $\beta_2 = 19$, and $\beta_3 = 28$ (see [13]). Initial temperature was $T_0 = 2$ and was decreased according to the formula $T_k = T_0 \, \text{ld}(k/4)$ with k being the iteration number. All results shown in Fig. 4 were obtained after 600 iterations.

This initial solution of the surface orientation, n, was correct at the occluding boundary, and $n = [0, 0, 1]$ elsewhere. Surface albedo, $\rho(x, y)$, was initialized to a random value for the first iteration and was allowed to vary in steps of 0.2. Discontinuity line elements were initialized randomly (with probability $p = 0.5$) to "on" or "off".

As can be seen from Fig. 4, recovery of surface orientation and albedo was perfect except in a band of varying width near the occluding boundary. This effect was present with a large range of other parameter values. The reasons for this will be discussed below.

Figure 5 shows average errors of the inferred surface orientation, $E[n_{\text{inferred}} \cdot n_{\text{true}}]$, and average errors of the inferred surface albedo, $E[\rho_{\text{inferred}} - \rho_{\text{true}}]$ for the egg-shaped Mondrian surface over iterations. They decrease exponentially, reaching an asymptote after about 200 iterations. The final errors are due to an imperfect fit near the occluding boundary, whereas there is virtually no error in shape/albedo recovery in the interior of the regions.

4. Discussion

The results show that our proposed method is capable of simultaneously recovering surface orientation and surface albedo for surfaces with piecewise constant albedo using shading information. This is achieved with a much higher computational effort (the results presented were obtained after 600 iterations) than other known methods (about 50 iterations for the Brooks and Horn method and a single iteration for Pentland's method).

It is not immediately obvious why we did not use a unified approach both for recovering albedo and for recovering surface orientation. Stochastic relaxation is computationally feasible only if each node can assume a small, fixed number of states, as we have assumed for albedo. Surface orientation, on the other hand, varies continuously and thus cannot be estimated efficiently using simulated annealing. Hybrid approaches similar to ours have also been suggested for recovering piecewise continuous surfaces from noisy data (Marroquin *et al.* 1987; Hutchinson and Koch 1986).

The poor performance in recovering surface orientation and albedo near the occluding boundary is caused by the smoothing term in [3]. In its discrete form it produces surfaces that are too flat and it does not enforce smoothness everywhere, as can be seen from the fact that the smoothness term alone (i.e., using $\alpha \rightarrow \infty$ in [3]) produces a conelike shape if the initial solution is correct at the occluding boundary. In our case this leads to a consistent overestimation of albedo, $\hat{\rho}(x, y)$, near the occluding boundaries, as can be seen in Fig. 4. This deficiency of the smoothing term was not apparent in Brooks and Horn (1985), as it was compensated by the first term in [3], the irradiance prediction term.

One possible way to overcome this problem is to use a higher-order smoothing term in [3] such as, for example, the biharmonic operator (Terzopoulos 1984; Horn and Brooks 1986). Although the biharmonic operator does enforce smoothness, it leads, on the other hand, to new problems, such as more complex boundary conditions (Horn and Brooks 1986, p. 186f) and a much slower convergence rate in the iterative estimation (Brandt 1977). The latter problem can be solved by resorting to multi-grid relaxation in the iterative estimation. Work is currently under progress to study the feasibility of this approach.

In the section on combining the recovery of surface orientation with the recovery of surface albedo (Sect. 2.3), we discussed three different approaches to combining the two mechanisms, eventually favouring the simultaneous recovery approach. We argued that the albedo-surface approach in which boundaries of "albedo patches" are localized first using some edge detection mechanism was not reliable enough for identifying regions of different albedo. However, this information can be used to initialize and (or) to constrain the introduction of discontinuity line elements in the recovery of the albedo map. The latter can be realized by introducing a data term similar to [10] for the line process. We expect an increase in efficiency of the recovery process using this information and are currently extending our method along these lines.

Acknowledgements

This research was partially supported by grant 81.166.0.84 of the Swiss National Science Foundation and by grant OGP0038521 of the Natural Sciences and Engineering Research Council of Canada to the first author and partially supported by the Consiglio Nazionale delle Ricerche to the second author. The authors would like to thank the two anonymous reviewers for their very helpful comments.

BLAKE, A., and ZISSERMAN, A. 1987. Visual reconstruction. MIT Press, Cambridge, MA.

BRANDT, A. 1977. Multi-level adaptive solutions to boundary value problems. Mathematical computation, **31**: 333–390.

BRELSTAFF, G., and BLAKE, A. 1980. Detecting specular reflections using Lambertian constraints. Proceedings of the Second International Conference on Computer Vision, Tampa, FL, pp. 297–302.

BROOKS, M.J., and HORN, B.K.P. 1985. Shape and source from shading. Proceedings of the 9th International Joint Conference on Artificial Intelligence, Los Angeles, CA, pp. 932–936.

FRANKOT, R.T., and CHELLAPPA, R. 1987. A method for enforcing integrability in shape from shading algorithms. Proceedings of the First International Conference on Computer Vision, London, England, June 8–11, pp. 118–127.

GEMAN, S., and GEMAN, D. 1984. Stochastic relaxation, Gibbs distribution, and the Bayesian restoration of images. IEEE Transactions on Pattern Analysis and Machine Intelligence, **PAMI-6**: 721–741.

HADAMARD, J. 1923. Lectures on the Cauchy problem in linear partial differential equations. Yale University Press, New Haven, CT.

HEALEY, G., and BINFORD, T.O. 1986. Local shape from specularity. Report No. STAN-CS-86-1139, Department of Computer Science, Stanford University, Stanford, CA.

HORN, B.K.P., and BROOKS, M.J. 1986. The variational approach to shape from shading. Computer Vision, Graphics, and Image Processing, **33**: 174–208.

HUTCHINSON, J.M., and KOCH, C. 1986. Simple analog and hydrid networks for surface interpolation. Proceedings of the AIP Conference 151 on Neural Networks for Computing, Snowbird, UT. American Institute of Physics, New York, NY. pp. 235–240.

HUTCHINSON, J.M., KOCH, C., LUO, J., and MEAD, C. 1988. Computing motion using analog and binary resistive networks. IEEE Computer Magazine, **21**(March): 52–64.

KIRKPATRICK, S., GELATT, C.D., and VECCHI, M.P. 1983. Optimization by simulated annealing. Science, **220**: 671–680.

MARROQUIN, J., MITTER, S., and POGGIO, T. 1987. Probabilistic solution of ill-posed problems in computational vision. Journal of the American Statistical Association, **82**(397): 76–89.

METROPOLIS, N., ROSENBLUTH, A., ROSENBLUTH, M., TELLER, A., and TELLER, E. 1953. Equation of state calculations by fast computing machines. Journal of Physical Chemistry, **21**: 1087.

PAPOULIAS, A. 1985. Curve segmentation using weak continuity constraints. M.Sc. thesis, Department of Computer Science, University of Edinburgh, Edinburgh, Scotland.

PENTLAND, A. 1984. Local shading analysis. IEEE Transactions on Pattern Analysis and Machine Intelligence, **PAMI-6**: 170–187.

POGGIO, T., GAMBLE, E.B., and LITTLE, J.J. 1988. Parallel integration of vision modules. Science, 242 (4877): 436–440.

TERZOPOULOS, D. 1985. Integrating visual information for multiple sources for the cooperative computation of surface shape. *In* Pixels to predicates: recent advances in computational and robotic vision. *Edited by* A. Pentland. Ablex Publishing Company, Norwood, NJ.

TIKHONOV, A.N. and ARSENIN, V.Y. 1977. Solutions of ill-posed problems. Winston & Sons, Washington, DC.

TORRE, V., and POGGIO, T.A. 1986. On edge detection. IEEE Transactions on Pattern Analysis and Machine Intelligence, **PAMI-8**: 147–163.