

The Computation of Visible-Surface Representations

DEMETRI TERZOPOULOS, MEMBER, IEEE

Abstract—The low-level interpretation of images provides constraints on 3-D surface shape at multiple resolutions, but typically only at scattered locations over the visual field. Sparse constraints from many sources collect into visible-surface representations where, as a precursor to higher-level visual tasks, intermediate-level processing reconstructs multiscale surface shape information at every image point. This paper develops a computational theory of visible-surface representations. The visible-surface reconstruction process that computes these quantitative representations unifies formal solutions to the key problems of 1) integrating multiscale constraints on surface depth and orientation from multiple visual sources, 2) interpolating dense, piecewise smooth surfaces from these constraints, 3) detecting surface depth and orientation discontinuities to impose boundary conditions on interpolation, and 4) structuring large-scale, distributed surface representations to achieve computational efficiency. Visible-surface reconstruction is an inverse problem. A well-posed variational formulation results from the use of a controlled-continuity surface model. Discontinuity detection amounts to the identification of this generic model's distributed parameters from the data. Finite element shape primitives yield a local discretization of the variational principle. The result is an efficient algorithm for visible-surface reconstruction. The algorithm deploys numerical relaxation in multigrid hierarchies and is suited to implementation on massively parallel networks of locally interconnected processors. Several applications contribute to an empirical evaluation of the framework.

Index Terms—Discontinuity detection, finite elements, multigrid relaxation, multiresolution methods, multisource integration, piecewise continuous reconstruction, regularization, variational principles, visible-surface representations.

I. INTRODUCTION

OVER 30 years ago, J. J. Gibson made a seminal conjecture: human visual perception in the natural environment amounts to perception of *visible surfaces* [26]. The representation of visible surfaces has since attracted considerable interest as an intermediate goal of computational vision. A variety of low-level visual processes participate in the recovery of the 3-D information from 2-D image data. The information contributed by each process partially constrains the shapes of visible surfaces. These constraints collect into *visible-surface representations* where intermediate-level processing takes place. The result is a dense, explicit description of the shapes and

configurations of visible surfaces. Visible-surface representations may be the basis of surface perception and they can provide quantitative information vital to higher-level surface analysis and object recognition tasks [40]. The computational vision literature describes several visible-surface representations, including depth and needle maps [34], intrinsic images [2], $2\frac{1}{2}$ -D sketches [41], and multiresolution representations [51].

This paper develops a computational approach to intermediate-level vision. The *visible-surface reconstruction process* proposed for generating and dynamically updating visible-surface representations, a generalization of the distributed, data-driven algorithm developed in [51], unifies the following four computational goals [52].

1) *Integration*: The visible-surface reconstruction process integrates local surface shape constraints from multiple sources and it fuses this information across multiple scales of resolution. The human visual system provides evidence for the importance of integration. It copes with a broad range of spatial structure in natural scenes by generating multiscale surface shape constraints through multiple spatial frequency channels [11]. Moreover, it coordinates two categories of low-level shape estimation processes [40]: The first, "correspondence" processes such as stereopsis and structure-from-motion, typically involves multiple image frames separated by relatively small space-time intervals. Correspondence processes triangulate interframe spatiotemporal disparities between corresponding surface features to derive *depth constraints*—estimates of range from the viewer to positions on visible surfaces. The second category comprises the "shape-from" processes, which can operate on single image frames. By accounting for the projective distortion of imaged surface properties such as shading, texture, and bounding contours, these processes derive *orientation constraints*—estimates of local surface attitude relative to the viewer. The visible-surface reconstruction process resolves ambiguities, and it counteracts the detrimental effects of noise and inaccuracies by integrating multiresolution depth and orientation constraints.

2) *Interpolation*: The visible-surface reconstruction process continuously propagates the integrated shape information into regions lacking shape constraints. Image representations make explicit certain local features (edges, markings, texture boundaries, etc.) correlated to salient events on physical surfaces. Since these features do not occur densely over the visual field, constraints generated by the low-level shape estimation processes will also be scattered over a subset of image points. The human visual

Manuscript received June 7, 1985; revised December 23, 1987. Recommended for acceptance by W. B. Thompson. This paper is based on MIT AI Memo 800, March 1985, by the author. Support for the Artificial Intelligence Laboratory, Massachusetts Institute of Technology, was provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research Contract N00014-75-C-0643 and the System Development Foundation. The author was supported by the Natural Sciences and Engineering Research Council of Canada and the Fonds F.C.A.C., Quebec, Canada.

The author is with Schlumberger Palo Alto Research, 3340 Hillview Avenue, Palo Alto, CA 94304.

IEEE Log Number 8821771.




Fig. 1. A sparse random dot stereogram. Binocular fusion elicits a percept of dense planar surfaces. A central, opaque, textured surface is perceived suspended nearer in depth over a similarly textured background. Vivid depth discontinuities separate the dense surfaces.

system, however, systematically interprets sparse visual stimuli, such as random dot stereograms (see Fig. 1), as dense, coherent 3-D surfaces, even when the dot density is reduced so that depth is indeterminate over as much as 98 percent of the visible surface area (see, e.g., the psychophysical study [18]). The interpolatory action of the visible-surface reconstruction process accounts for this phenomenon of “*filling in the gaps.*”

3) *Discontinuities*: The visible-surface reconstruction process collects into discontinuity maps all low-level information about surface discontinuities, including intensity, texture, and motion boundary fragments. The process autonomously refines these maps, detecting and localizing additional surface discontinuities. Perceptually crucial are *depth discontinuities*, typically contours along which a surface in the scene occludes itself or another surface, as well as *orientation discontinuities* along creases or cusps of a continuous surface. Accordingly, random dot stereograms not only give impressions of coherent surfaces, but they also elicit vivid percepts of surface discontinuities at abrupt disparity changes (see Fig. 1). The refined discontinuity maps computed by the visible-surface reconstruction process provide (dynamic) boundary conditions which limit the interpolation of shape constraints.

4) *Efficiency*: The visible-surface reconstruction process must efficiently produce large quantities of numerical shape information. Visible-surface representations form and evolve in real time, in spite of the immense computational burden associated with surface reconstruction at foveal resolution. Massive, fine-grained parallelism appears to be the most viable computational architecture for this purpose. However, a characteristic limitation of massively parallel hardware is the lack of high bandwidth connections except between neighboring processors. The consequent propagation delays, however, can severely hamper global exchange of information across large visual representations. The visible-surface reconstruction process overcomes this potential inefficiency by employing *multiresolution relaxation* within a hierarchy of surface representations each attuned to a range of spatial scales [51], [55].

II. MATHEMATICAL BASIS OF VISIBLE-SURFACE RECONSTRUCTION

Let the true distance from the viewer to visible surfaces be represented as $z = Z(x, y)$, a function of the image coordinates x and y . Low-level visual processes generate

a set of noise corrupted surface shape estimates (i.e., constraints) $\{c_i\}$ which can be expressed using the notation

$$c_i = \mathcal{L}_i[Z(x, y)] + \epsilon_i, \quad (1)$$

where \mathcal{L}_i denotes measurement functionals of $Z(x, y)$ and ϵ_i denotes associated measurement errors. An abstract statement of the visible-surface reconstruction problem is: reconstruct from the available constraints $\{c_i\}$ the depth function $Z(x, y)$ along with an explicit representation of its discontinuities over the visual field.

A. Regularizing the Inverse Problem

Visible-surface reconstruction is a nontrivial inverse problem. First, coincident but slightly inconsistent shape estimates from different visual processes will locally overdetermine surface shape. Second, sparse constraints scattered over the visual field restrict surface shape locally, but do not determine it uniquely everywhere; there remain infinitely many feasible surfaces. Third, shape estimates are subject to errors, and high spatial frequency additive noise, regardless how small its (RMS) amplitude, can locally perturb the surface (orientation) radically.

The inverse problem is ill-posed because the above three considerations preclude any prior guarantee that the solution will exist, or that it will be unique, or that it will be stable with respect to measurement errors. The regularization method [57], [44], [56] provides a systematic approach to reformulating this ill-posed inverse problem as a well-posed and effectively solvable variational principle.

Let \mathcal{J} be a linear space of admissible functions. Let $\mathcal{S}(v)$ be a stabilizing functional which measures the (lack of) smoothness of a function $v \in \mathcal{J}$. Let $\mathcal{P}(v)$ be a penalty functional on \mathcal{J} which provides a measure of the discrepancy between v and the given constraints. The regularized visible-surface reconstruction problem is formulated according to the following variational principle [51]:

VPI: Find $u \in \mathcal{J}$ such that

$$\mathcal{E}(u) = \inf_{v \in \mathcal{J}} \mathcal{E}(v), \quad (2)$$

where the energy functional

$$\mathcal{E}(v) = \mathcal{S}(v) + \mathcal{P}(v). \quad (3)$$

The solution $u(x, y)$ characterizes the best reconstruction of the depth function $Z(x, y)$ as the smoothest admissible function $v \in \mathcal{J}$ which is most compatible with the available constraints. When the solution exists, it satisfies the Euler-Lagrange equations which express the necessary condition for the minimum as the vanishing of the first variational derivative δ_u of the energy functional:

$$\delta_u \mathcal{E}(u) = \delta_u \mathcal{S}(u) + \delta_u \mathcal{P}(u) = 0. \quad (4)$$

The next two sections examine $\mathcal{S}(v)$ and $\mathcal{P}(v)$.

B. A Controlled-Continuity Surface Model

To accomplish the reconstruction, the stabilizing functional imposes generic continuity conditions on functions

admissible as possible solutions. Such conditions are tenable inasmuch as the coherence of matter tends to give rise to continuous or smooth surfaces relative to the viewing distance over some spatial resolution range.

Controlled-continuity stabilizers which provide local control over the continuity of the solution enable the problem to be regularized while preserving surface discontinuities. The controlled-continuity stabilizer of order 2 in two dimensions suffices in reconstructing nominally C^1 continuous surfaces (continuously varying surface normal) along with explicit depth and orientation discontinuities. The stabilizer is given by

$$\mathcal{S}_{\rho\tau}(v) = \frac{1}{2} \int_{\Omega} \int_{\Omega} \rho(x, y) \left\{ \tau(x, y) (v_{xx}^2 + 2v_{xy}^2 + v_{yy}^2) + [1 - \tau(x, y)] (v_x^2 + v_y^2) \right\} dx dy, \quad (5)$$

where $\Omega \subset \mathbb{R}^2$ denotes the image domain, and $\rho(x, y)$ and $\tau(x, y)$ are real-valued weighting functions whose range is $[0, 1]$. Assuming natural (i.e., free) boundary conditions on $\partial\Omega$, the variational derivative of (5) in the interior of Ω is given by

$$\begin{aligned} \delta_v \mathcal{S}_{\rho\tau}(v) &= \frac{\partial^2}{\partial x^2} (\mu v_{xx}) + 2 \frac{\partial^2}{\partial x \partial y} (\mu v_{xy}) + \frac{\partial^2}{\partial y^2} (\mu v_{yy}) \\ &\quad - \frac{\partial}{\partial x} (\eta v_x) - \frac{\partial}{\partial y} (\eta v_y), \end{aligned} \quad (6)$$

where $\mu(x, y) = \rho(x, y)\tau(x, y)$ and $\eta(x, y) = \rho(x, y)[1 - \tau(x, y)]$.

With $\mathcal{S}(v) \equiv \mathcal{S}_{\rho\tau}(v)$ in (3), the values of the *continuity control functions* $\rho(x, y)$ and $\tau(x, y)$ at any point $(x, y) \in \Omega$ determine the local continuity of $u(x, y)$ at that point: $\lim_{\tau(x,y) \rightarrow 0} \mathcal{S}_{\rho\tau}(v)$ locally characterizes a membrane spline, a C^0 surface that need only be continuous; $\lim_{\tau(x,y) \rightarrow 1} \mathcal{S}_{\rho\tau}(v)$ locally characterizes a thin-plate spline, a C^1 surface which is continuous and has continuous first derivatives; $\lim_{\rho(x,y) \rightarrow 0} \mathcal{S}_{\rho\tau}(v)$ characterizes a locally discontinuous surface.¹ Intermediate values of $\rho(x, y)$ and $\tau(x, y)$ locally characterize a hybrid C^1 “thin-plate surface under tension,” where $\rho(z, y)$ is a spatially varying surface “cohesion” and $[1 - \tau(x, y)]$ is the spatially varying surface “tension” [53], [56].

Hence, the continuity control functions $\rho(x, y)$ and $\tau(x, y)$ constitute an explicit representation of depth and orientation discontinuities, respectively, over the visual field Ω . A subsequent section examines the automatic identification of these functions to estimate discontinuities unknown *a priori* but implicit in the data $\{c_i\}$.

C. A Penalty Functional

A reasonable penalty functional for (3) is the weighted Euclidean norm

$$\mathcal{P}(v) = \frac{1}{2} \sum_i \alpha_i (\mathcal{L}_i[v] - c_i)^2, \quad (7)$$

¹To have the intended effect on the continuity of the solution $u(x, y)$ to *VPI*, $\rho(x, y)$ and $\tau(x, y)$ must vanish on a set of nonzero measure. In the discrete formulation of the problem (see below) the finite element surface primitives automatically provide the necessary finite support.

where the α_i are nonnegative real-valued constraint parameters. For $\alpha_i = 1/\lambda\sigma_i^2$ with λ a proportionality factor, this functional is in fact optimal for independently distributed measurement errors ϵ_i in (1) with zero means and variances σ_i^2 .

Measurement functionals for surface reconstruction may be synthesized from point evaluation of generalized k th-order derivatives: $\mathcal{L}_i[v] = (\partial^k v / \partial x^j \partial y^{k-j})|_{(x_i, y_i) \in \Omega}$, for $j = 0, 1, \dots, k$. Hence, zeroth-order (evaluation) functionals $\mathcal{L}_i[v(x, y)] = v(x_i, y_i)$ serve to model the set of local depth constraints $c_i = v(x_i, y_i) + \epsilon_i = d_{(x_i, y_i)}$, for $i \in D$. The local surface orientation, determined by the components of the surface normal $\mathbf{n}(x_i, y_i) = [v_x(x_i, y_i), v_y(x_i, y_i), -1]$, is handled in turn by the two first-order derivative functionals $\mathcal{L}_i[v(x, y)] = v_x(x_i, y_i)$ and $\mathcal{L}_i[v(x, y)] = v_y(x_i, y_i)$. This yields analogous expressions for the local orientation constraints: the set $c_i = v_x(x_i, y_i) + \epsilon_i = p_{(x_i, y_i)}$, for $i \in P$, and the set $c_i = v_y(x_i, y_i) + \epsilon_i = q_{(x_i, y_i)}$, for $i \in Q$. It is straightforward to synthesize additional functionals—e.g., involving directional or higher-order derivatives (for curvature constraints, etc.).

The penalty functional used in the sequel is written as

$$\begin{aligned} \mathcal{P}(v) &= \frac{1}{2} \sum_{i \in D} \alpha_{d_i} [v(x_i, y_i) - d_{(x_i, y_i)}]^2 \\ &\quad + \frac{1}{2} \sum_{i \in P} \alpha_{p_i} [v_x(x_i, y_i) - p_{(x_i, y_i)}]^2 \\ &\quad + \frac{1}{2} \sum_{i \in Q} \alpha_{q_i} [v_y(x_i, y_i) - q_{(x_i, y_i)}]^2, \end{aligned} \quad (8)$$

where the α_i parameters are now distinguished as α_{d_i} , α_{p_i} , and α_{q_i} .

D. Physical Interpretation

A physical model of variational principle *VPI* is illustrated in Fig. 2. The controlled-continuity stabilizer models an elastic surface whose energy of deformation $\mathcal{S}_{\rho\tau}(v)$ compels its shape to vary smoothly almost everywhere (but not at discontinuities). Constraints apply forces in the z direction which deflect the surface from its nominally planar state; the penalty functional $\mathcal{P}(v)$ is the total deformation energy of a set of ideal springs attached to the constraints. The infrastructure of scattered depth constraints determine the deflection $u(x, y)$ of the elastic surface at equilibrium, as illustrated by Fig. 2(a). The height of the constraint encodes the magnitude of the local depth estimate. The tightness of each constraint is controlled by the associated spring stiffness α_{d_i} . Fig. 2(b) illustrates an orientation constraint coercing the local surface normal; the constraint parameters α_{p_i} and α_{q_i} control the spring stiffness.

E. Existence, Uniqueness, and Stability of the Solution

Existence, uniqueness, and stability of the solution $u(x, y)$ to *VPI* are guaranteed when the functional $\mathcal{E}_{\rho\tau}(v) = \mathcal{S}_{\rho\tau}(v) + \mathcal{P}(v)$ is a norm in the admissible space \mathcal{H} . Now, $\mathcal{E}_{\rho\tau}(v)$ is *a priori* only a seminorm in \mathcal{H} , (a particular class of Sobolev spaces); the null space \mathcal{N} of func-




Fig. 2. The physical model. (a) Thin-plate surface under tension and depth constraints. (b) Local influence of an orientation constraint.

tions v that it maps to zero is at worst (for the thin-plate case, $\tau = 0$) the space of all polynomials over \mathbb{R}^2 of degree one or less [56]. The penalty functional $\mathcal{P}(v)$ will make $\mathcal{E}_{\rho\tau}(v)$ a norm, however, provided it at least constrains \mathcal{U} to a unique linear polynomial. This will be the case if \mathcal{L}_i includes evaluation functionals at 3 points defining a unique linear polynomial (an \mathcal{U} -unisolvant set). It is possible to prove the following [53]:

The solution $u(x, y)$ to VPI will exist, be unique, and stable given at least

- 1) three noncolinear depth constraints, or
- 2) two depth constraints and a single p or q constraint, or
- 3) a single depth constraint and a single p and q constraint, or
- 4) a single p and q constraint with the "center of gravity" of the surface fixed.

At least one of these conditions is satisfied in practice, due to the large number of constraints typically provided by the low-level shape estimation processes (the center of gravity can be constrained when necessary). Thus the visible-surface reconstruction problem VPI is well-posed in practice: hence, for preset $\rho(x, y)$ and $\tau(x, y)$ the plate/spring system possesses a unique state of stable equilibrium—the minimal energy state $u(x, y)$.

III. IDENTIFICATION OF SURFACE DISCONTINUITIES

It is natural to view the problem of detecting surface discontinuities as one of distributed parameter identification within a variational formulation of visible-surface reconstruction [56] (distributed parameter identification is a problem of widespread interest [45], [16]). The parameter functions to be identified (estimated from the available data) are the continuity control functions $\rho(x, y)$ and $\tau(x, y)$ in the stabilizer (5). The present section proposes two estimation procedures which dynamically adjust the controlled-continuity model during surface reconstruction such that its continuity becomes consistent with discontinuities implied by the data. The first procedure detects discontinuities by locally monitoring sharp deflections in the evolving surface. The second extends variational principle VPI to govern the estimation of discontinuities according to generic shape constraints.

A. Discontinuity Identification by Local Validation

Fig. 3 illustrates in cross section a C^1 continuous portion of reconstructed surface attempting to interpolate over

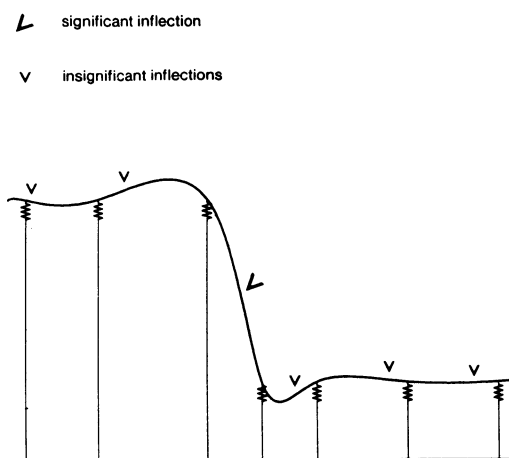


Fig. 3. Cross-section of reconstructed surface across a depth discontinuity indicating significant and insignificant surface inflections.

an implicit depth discontinuity. The constraints on either side of the discontinuity impart opposing bending moments, thus causing a surface inflection indicated by a sign change in the total bending moment. This can form the basis of a validation test for determining those neighborhoods of a solution where the continuity assumption is invalid. Note also how the surface overshoots the constraints because its smoothness opposes the sudden jump in depth. The spurious inflections are relatively weak and they may be rejected using a significance measure associated with the validation test. A technically suitable measure having some psychophysical justification may be based on the depth gradient at an inflection.

For a thin-plate deflection function $u(x, y)$, the bending moment per unit length parallel to the $x - z$ plane is proportional to $-u_{xx}$, while its counterpart parallel to the $y - z$ plane is proportional to $-u_{yy}$ [50]. The total moment is $M(x, y) = -(u_{xx} + u_{yy}) = -\Delta u$, where Δ denotes the Laplacian operator. A zero-crossing in M may be interpreted as a significant depth discontinuity if $G(x, y) = |\nabla u|^2 = u_x^2 + u_y^2$ exceeds a limit t_d^2 . The limit t_d should be large enough to reject weak inflections, while not so conservative as to miss many true depth discontinuities. A possible criterion for choosing t_d is suggested by Panum's limiting case; i.e., when a tilted surface begins to occlude itself from one eye, causing stereopsis to fail. A tighter criterion is the roughly isotropic disparity gradient limit between fusion and diplopia of approximately 1 (only half the Panum limit), as measured by Burt and Julesz [15].

An analogous local validation test is devised for orientation discontinuities. Relative extrema of the bending moment (local extrema of curvature) localize orientation discontinuities. The associated significance criterion requires the magnitude of the total bending moment $|M(x, y)|$ of the surface to exceed a (high curvature) limit t_o . The sign of a bending moment extremum indicates the sense of the orientation discontinuity; negative indicates a concave crease, and positive, a convex crease.

The local validation procedure for reconstructing piecewise continuous surfaces is as follows.

1) Reconstruct a tentative C^1 continuous surface on Ω ; i.e., solve *VPI* with $\rho(x, y) = \tau(x, y) = 1$.

2) Introduce significant depth discontinuities into the resulting surface; i.e., set $\rho(x, y) = 0$ at $\{(x, y) \mid M(x, y) = 0 \text{ and } G(x, y) > t_d\}$ and continue the reconstruction by solving *VPI*.

3) Introduce significant orientation discontinuities into the resulting surface; i.e., set $\tau(x, y) = 0$ at $\{(x, y) \mid \nabla M(x, y) = 0 \text{ and } |M(x, y)| > t_o\}$ and continue the reconstruction by solving *VPI*.

4) Repeat steps 2) and 3) with decreasing t_d and t_o .

Step 4) sets up an iterative *continuation* cycle for solving *VPI*, where steps 2) and 3) continue using the surface resulting from the immediately preceding step (here, (4) becomes a quasilinear equation due to the dependence of ρ and τ not only on position but also on partial derivatives of u). It is wasteful to compute each approximation to high accuracy, since it serves only as an initial condition toward computing a better approximation over the succeeding cycle.

The local validation procedure is reminiscent of the common practice of detecting intensity edges in image functions by applying thresholded local difference operations. Since a local edge operator, such as a Laplacian, is easily corrupted by noise, a smoothing prefilter is usually applied to the image to improve the response. $\mathcal{S}_{\rho\tau}$ has an analogous smoothing effect on scattered, noisy shape constraints (standard low-pass filters, such as Gaussians, are inapplicable to irregular samples). While regularization based smoothing permits the reliable computation of numerical derivatives in continuous regions [58], the smoothing property of the tentative surface computed in step 1) tends to obscure subtle discontinuities [56]. This problem is also typical of smoothing edge detection operators [39].

B. Discontinuity Identification by Variational Continuity Control

The second discontinuity identification approach embeds *VPI* within an outer variational principle which estimates the continuity control parameters $\rho(x, y)$ and $\tau(x, y)$. The surface is permitted to crease and fracture in order to reduce the total energy below the minimum obtainable with a globally C^1 surface. This resolves the conflict between discontinuity identification and regularization smoothing.

Variational continuity control is formulated in terms of the following variational problem:

VP2: Find u , ρ^* , and τ^* such that

$$\mathcal{E}(u, \rho^*, \tau^*) = \inf_{\rho, \tau} \mathcal{E}(u, \rho, \tau), \quad (9)$$

where

$$\mathcal{E}(u, \rho, \tau) = \mathcal{D}(\rho, \tau) + \inf_v \mathcal{E}_{\rho\tau}(v); \quad (10)$$

$$\mathcal{E}_{\rho\tau}(v) = \mathcal{S}_{\rho\tau}(v) + \mathcal{P}(v). \quad (11)$$

Assuming suitable continuity, $u(x, y)$, $\rho^*(x, y)$, and $\tau^*(x, y)$ satisfy the coupled, nonlinear Euler-Lagrange

equations

$$\begin{aligned} \delta_u \mathcal{E}(u, \rho^*, \tau^*) &= \frac{\partial^2}{\partial x^2} (\mu u_{xx}) + 2 \frac{\partial^2}{\partial x \partial y} (\mu u_{xy}) \\ &+ \frac{\partial^2}{\partial y^2} (\mu u_{yy}) - \frac{\partial}{\partial x} (\eta u_x) \\ &- \frac{\partial}{\partial y} (\eta u_y) + \delta_u \mathcal{P}(u) = 0; \end{aligned}$$

$$\begin{aligned} \delta_{\rho^*} \mathcal{E}(u, \rho^*, \tau^*) &= \tau^* (u_{xx}^2 + 2u_{xy}^2 + u_{yy}^2) \\ &+ [1 - \tau^*] (u_x^2 + u_y^2) \\ &+ \delta_{\rho^*} \mathcal{D}(\rho^*, \tau^*) = 0; \end{aligned}$$

$$\begin{aligned} \delta_{\tau^*} \mathcal{E}(u, \rho^*, \tau^*) &= \rho^* [(u_{xx}^2 + 2u_{xy}^2 + u_{yy}^2) \\ &- (u_x^2 + u_y^2)] + \delta_{\tau^*} \mathcal{D}(\rho^*, \tau^*) = 0, \end{aligned} \quad (12)$$

with suitable boundary conditions on $\partial\Omega$.

The functional $\mathcal{D}(\rho, \tau)$ maps $\rho(x, y)$ and $\tau(x, y)$ into a nonnegative energy. Its role in the outer problem is analogous to that of $\mathcal{S}_{\rho\tau}(v)$ in the inner problem *VPI*: it serves as a stabilizer for estimating the continuity control functions from the available surface shape constraints. Locally reducing the smoothness of the surface will reduce its resistance to sudden deflections imposed by the data and hence release potential energy; i.e., from (5), $\mathcal{S}_{\rho\tau}(v)$ considered as a function of (v, ρ, τ) decreases as $\int \int_{\Omega} \rho(x, y) dx dy$ and $\int \int_{\Omega} \tau(x, y) dx dy$ decrease. The introduction of discontinuities must be penalized, however, because $\rho(x, y) = 0$ everywhere would trivially minimize the energy. The penalty can simply increase monotonically according to a total discontinuity measure; e.g.,

$$\begin{aligned} \mathcal{D}(\rho, \tau) &= \int \int_{\Omega} \beta_d [1 - \rho(x, y)] \\ &+ \beta_o [1 - \tau(x, y)] dx dy, \end{aligned} \quad (13)$$

where β_d and β_o are positive energy scaling parameters for the depth and orientation discontinuity contributions, respectively.

More interestingly, $\mathcal{D}(\rho, \tau)$ can express a predisposition for discontinuities arranged along piecewise continuous contours in the $x - y$ plane (making allowances with regard to the condition in footnote 1). An appealing formulation is in terms of curvilinear controlled-continuity constraints; for instance,

$$\oint_C \beta_b(s) \{ \beta_a(s) |c_{ss}|^2 + [1 - \beta_a(s)] |c_s|^2 \} ds, \quad (14)$$

the curvilinear analog of $\mathcal{S}_{\rho\tau}$, where s denotes arc length along contours $c(s) = [x(s), y(s)]$ in C , a collection of discontinuity contours. Here, locally zeroing β_b allows break points, while locally zeroing β_a allows tangent discontinuity points to form in the contours. Again, these events require energy penalties. Naturally, the embedded

structure of the variational continuity control problem reflects the recursive embedding of visual singularities—from surfaces to contours to points (see [56]).

Although the inner energy functional $\mathcal{E}_{\rho\tau}(v)$ has a unique minimum for fixed ρ and τ (it is quadratic, convex, and positive definite, given the conditions of Section II-E), this will certainly not be the case for $\mathcal{E}(v, \rho, \tau)$ in (10) which permits variation in the continuity control functions. Embedded variational problems such as VP2 are nonconvex typically and may be nondifferentiable (e.g., when the distributed parameters do not vary continuously; for instance, their range is the binary set $\{0, 1\}$). The rigorous treatment of such problems is the goal of nondifferentiable analysis and nonconvex optimization theory [46], [17], and is beyond the scope of the present paper. Although applicable, neither deterministic subgradient-type methods nor stochastic optimization techniques [21] have been attempted here. Instead, a discrete continuation approach is taken to solving the variational continuity control problem [61]. A subsequent section proposes an iterative procedure which efficiently computes good, although not necessarily optimal solutions.

IV. THE DISCRETE SURFACE RECONSTRUCTION PROBLEM

A closed-form solution to the variational principle for visible-surface reconstruction is infeasible due to the irregular occurrence of constraints and discontinuities. Consequently, the continuous problem is approximated by a discrete variational principle whose solution may be computed numerically. To this end, the finite element method [49], [63] is an attractive local approximation technique [51]. As piecewise shape primitives defined in the viewer-centered coordinate system of visible-surface representations, finite elements are computationally compatible with the local depth and orientation constraints provided by the various shape estimation processes and they readily accommodate the irregular occurrence of constraints and discontinuities.

Although it is possible to discretize this problem using a variety of finite elements, including irregularly shaped isoparametric elements, the discretization will follow a regular Cartesian sampling pattern typical of images. The domain $\Omega \subset \mathbb{R}^2$ (assumed rectangular without loss of generality) is tessellated into square element subdomains with sides of length h . Nodes are located at subdomain corners where they are shared by adjacent subdomains. This results in a uniform Cartesian array of nodes that is well suited to VLSI implementation. The size h of the elements is adjustable so as to yield a one-to-one mapping between nodes and image pixels, as well as a geometric progression of coarser mappings. The nodes $\{ih, jh\} \cap \Omega$ are indexed by (i, j) for $i = 1, \dots, N_x^h$ and $j = 1, \dots, N_y^h$, where N_x^h and N_y^h are the number of nodes along the x and y axes, respectively.

The nodal variable $v_{i,j}^h = v(ih, jh)$ denotes the unknown displacement, or depth, at node (i, j) . Taken together, the $N^h = N_x^h \times N_y^h$ nodal variables form the vector

$\mathbf{v}^h \in \mathbb{R}^{N^h}$, to be determined by solving the discrete variational principle. Similarly, nodal parameters $\rho_{i,j}^h = \rho(ih, jh)$ and $\tau_{i,j}^h = \tau(ih, jh)$ represent the continuity control functions, and it is natural to interpret them as together comprising a *discontinuity map* in registration with the reconstructed surface.

A. The Discrete Equations

The reconstructed surface is represented by an assembly of square (nonconforming) finite elements E [51], each of which is a six-point full quadratic interpolant [see (24)]. \mathbf{v}^h determines the local interpolants which explicitly represent depth and orientation everywhere over the surface. The element leads to the following $O(h^2)$ finite difference type formulas for the required partial derivatives at node (i, j) :

$$\begin{aligned} v_{xx}^h &= \frac{1}{h^2} (v_{i+1,j}^h - 2v_{i,j}^h + v_{i-1,j}^h); \\ v_{yy}^h &= \frac{1}{h^2} (v_{i,j+1}^h - 2v_{i,j}^h + v_{i,j-1}^h); \\ v_{xy}^h &= \frac{1}{h^2} (v_{i+1,j+1}^h - v_{i,j+1}^h - v_{i+1,j}^h + v_{i,j}^h); \\ v_x^h &= \frac{1}{h} (v_{i+1,j}^h - v_{i,j}^h); \\ v_y^h &= \frac{1}{h} (v_{i,j+1}^h - v_{i,j}^h). \end{aligned} \quad (15)$$

Substitution of (15) along with the piecewise constant approximation $\rho(x, y) = \rho_{i,j}^h$ and $\tau(x, y) = \tau_{i,j}^h$ into (5), and noting that the area of each element is h^2 , yields the discrete functional

$$\begin{aligned} \mathcal{S}_{\rho\tau}^h(\mathbf{v}^h) &= \frac{1}{2} \sum_{i,j} \rho_{i,j}^h \left\{ \frac{\tau_{i,j}^h}{h^2} [(v_{i+1,j}^h - 2v_{i,j}^h + v_{i-1,j}^h)^2 \right. \\ &\quad + 2(v_{i+1,j+1}^h - v_{i,j+1}^h - v_{i+1,j}^h + v_{i,j}^h)^2 \\ &\quad + (v_{i,j+1}^h - 2v_{i,j}^h + v_{i,j-1}^h)^2] \\ &\quad + [1 - \tau_{i,j}^h] [(v_{i+1,j}^h - v_{i,j}^h)^2 \\ &\quad \left. + (v_{i,j+1}^h - v_{i,j}^h)^2] \right\}. \end{aligned} \quad (16)$$

Assuming a one-to-one mapping between nodes and image pixels, the constraints coincide with nodes (i, j) of the grid (but not all nodes need be constrained—note, the general case of a constraint occurring arbitrarily within an element domain E may be handled easily). To obtain a discrete expression for $\mathcal{P}(v)$, collect the nodes at which the various constraints occur in three sets; the set $(i, j) \in D$ at which depth constraints $d_{i,j}^h$ occur, and the sets $(i, j) \in P$ and $(i, j) \in Q$ at which orientation constraints $p_{i,j}^h$ and $q_{i,j}^h$ occur. Using symmetric difference approximations the

discrete form of (8) becomes

$$\begin{aligned} \mathcal{P}^h(\mathbf{v}^h) &= \frac{1}{2} \sum_{(i,j) \in D} \alpha_{d_{i,j}}^h (v_{i,j}^h - d_{i,j}^h)^2 \\ &+ \frac{1}{2} \sum_{(i,j) \in P} \alpha_{p_{i,j}}^h \left(\frac{1}{2h} (v_{i+1,j}^h - v_{i-1,j}^h) - p_{i,j}^h \right)^2 \\ &+ \frac{1}{2} \sum_{(i,j) \in Q} \alpha_{q_{i,j}}^h \left(\frac{1}{2h} (v_{i,j+1}^h - v_{i,j-1}^h) - q_{i,j}^h \right)^2. \end{aligned} \quad (17)$$

The energy-minimizing vector of nodal displacements \mathbf{u}^h satisfies the equilibrium condition

$$\nabla \mathcal{E}_{\rho\tau}^h(\mathbf{u}^h) = \nabla \mathcal{S}_{\rho\tau}^h(\mathbf{u}^h) + \nabla \mathcal{P}^h(\mathbf{u}^h) = 0, \quad (18)$$

where ∇ is the gradient operator [cf. (4)]. This generally nonlinear system of simultaneous equations reduces to a linear system for fixed $\rho_{i,j}^h$ and $\tau_{i,j}^h$ (i.e., for preset discontinuities), because $\mathcal{E}_{\rho\tau}^h(\mathbf{u}^h)$ is a quadratic form in the $u_{i,j}^h$. The basic, discrete surface reconstruction problem amounts to solving these *nodal equations* for \mathbf{u}^h .

The nodal equation at an arbitrary node (i, j) is given by $(\partial \mathcal{S}_{\rho\tau}^h(\mathbf{u}^h) / \partial u_{i,j}^h) + (\partial \mathcal{P}^h(\mathbf{u}^h) / \partial u_{i,j}^h) = 0$. Letting

$$\mu_{i,j}^h = \rho_{i,j}^h \tau_{i,j}^h / h^2 \quad \text{and} \quad \eta_{i,j}^h = \rho_{i,j}^h (1 - \tau_{i,j}^h), \quad (19)$$

the partial derivatives are

$$\begin{aligned} \frac{\partial \mathcal{S}_{\rho\tau}^h(\mathbf{u}^h)}{\partial u_{i,j}^h} &= \left\{ (u_{i,j}^h - 2u_{i-1,j}^h + u_{i-2,j}^h) \mu_{i-1,j}^h \right. \\ &+ (-2u_{i+1,j}^h + 4u_{i,j}^h - 2u_{i-1,j}^h) \mu_{i,j}^h \\ &+ (u_{i+2,j}^h - 2u_{i+1,j}^h + u_{i,j}^h) \mu_{i+1,j}^h \\ &+ (2u_{i,j}^h - 2u_{i-1,j}^h - 2u_{i,j-1}^h + 2u_{i-1,j-1}^h) \mu_{i-1,j-1}^h \\ &+ (-2u_{i+1,j}^h + 2u_{i,j}^h + 2u_{i+1,j-1}^h - 2u_{i,j-1}^h) \mu_{i,j-1}^h \\ &+ (-2u_{i,j+1}^h + 2u_{i-1,j+1}^h + 2u_{i,j}^h - 2u_{i-1,j}^h) \mu_{i-1,j+1}^h \\ &+ (2u_{i+1,j+1}^h - 2u_{i,j+1}^h - 2u_{i+1,j}^h + 2u_{i,j}^h) \mu_{i,j+1}^h \\ &+ (u_{i,j}^h - 2u_{i,j-1}^h + u_{i,j-2}^h) \mu_{i,j-1}^h \\ &+ (-2u_{i,j+1}^h + 4u_{i,j}^h - 2u_{i,j-1}^h) \mu_{i,j}^h \\ &+ (u_{i,j+2}^h - 2u_{i,j+1}^h + u_{i,j}^h) \mu_{i,j+1}^h \left. \right\} \\ &+ \left\{ (u_{i,j}^h - u_{i-1,j}^h) \eta_{i-1,j}^h \right. \\ &+ (u_{i,j}^h - u_{i+1,j}^h) \eta_{i,j}^h \\ &+ (u_{i,j}^h - u_{i,j-1}^h) \eta_{i,j-1}^h \\ &+ (u_{i,j}^h - u_{i,j+1}^h) \eta_{i,j+1}^h \left. \right\} \end{aligned} \quad (20)$$

[a discrete version of (6)], and for $(i, j) \in D \cap P \cap Q$,

$$\begin{aligned} \frac{\partial \mathcal{P}^h(\mathbf{u}^h)}{\partial u_{i,j}^h} &= (\alpha_{d_{i,j}}^h u_{i,j}^h - \alpha_{d_{i,j}}^h d_{i,j}^h) \\ &+ \left(\frac{\alpha_{p_{i-1,j}}^h}{4h^2} (u_{i,j}^h - u_{i-2,j}^h) - \frac{\alpha_{p_{i-1,j}}^h}{2h} p_{i-1,j}^h \right) \\ &+ \left(\frac{\alpha_{p_{i+1,j}}^h}{4h^2} (u_{i,j}^h - u_{i+2,j}^h) + \frac{\alpha_{p_{i+1,j}}^h}{2h} p_{i+1,j}^h \right) \\ &+ \left(\frac{\alpha_{q_{i,j-1}}^h}{4h^2} (u_{i,j}^h - u_{i,j-2}^h) - \frac{\alpha_{q_{i,j-1}}^h}{2h} q_{i,j-1}^h \right) \\ &+ \left(\frac{\alpha_{q_{i,j+1}}^h}{4h^2} (u_{i,j}^h - u_{i,j+2}^h) + \frac{\alpha_{q_{i,j+1}}^h}{2h} q_{i,j+1}^h \right). \end{aligned} \quad (21)$$

B. Computational Molecules

In the discrete equations, $\rho_{i,j}^h$ and $\tau_{i,j}^h$ may generally assume values continuously in $[0, 1]$. Nonetheless, the binary case in which only the endpoints of this range are permissible values ($\{0, 1\}$ indicating {presence, absence} of discontinuities) suggests the following graphical interpretation of the nodal equations:

Each term of (20) and (21) in parentheses may be visualized as a basic *computational molecule*. Molecules consist of *atoms*, indicated by circles, arranged in the spatial grid pattern and containing coefficients of the associated nodal variable. Fig. 4(a) illustrates the ten *plate molecules* obtained from the terms of the first component of (20), while Fig. 4(b) shows the four *membrane molecules* obtained from the terms of the second component. A double circle indicates the node (i, j) central to the nodal equation. Similarly, the depth constraint term in (21) yields the *depth constraint molecule* shown in Fig. 5(a). Associated with the latter is the factor $\alpha_{d_{i,j}}^h d_{i,j}^h$. The remaining orientation constraint terms of (21) yield the *orientation constraint molecules* and associated factors shown in Fig. 5(b).

Now visualize the formation of nodal equations as *molecular summation*, where the basic molecules combine at the central node (i, j) with coincident atoms summing together. Discontinuities, however, require *molecular inhibition*: if (i, j) is a discontinuity, $\mu_{i,j}^h$ or $\eta_{i,j}^h$ or both are zero, which inhibits the summation of certain molecules. Specifically:

1) Plate, depth constraint, and orientation constraint molecules sum at nondiscontinuity nodes.

2) Membrane and depth constraint molecules sum at orientation discontinuity nodes.

3) Orientation discontinuities inhibit plate and orientation constraint molecules.

4) Depth discontinuities inhibit all basic molecules.

At an interior node (i, j) , summation leads to the composite *nodal molecules* illustrated in Fig. 6. If node (i, j) is also a depth constraint, the depth constraint molecule and associated constraint factor sum with the nodal molecule for equation (i, j) shown in Fig. 6(a). Similarly,




Fig. 4. (a) Plate molecules. (b) Membrane molecules.

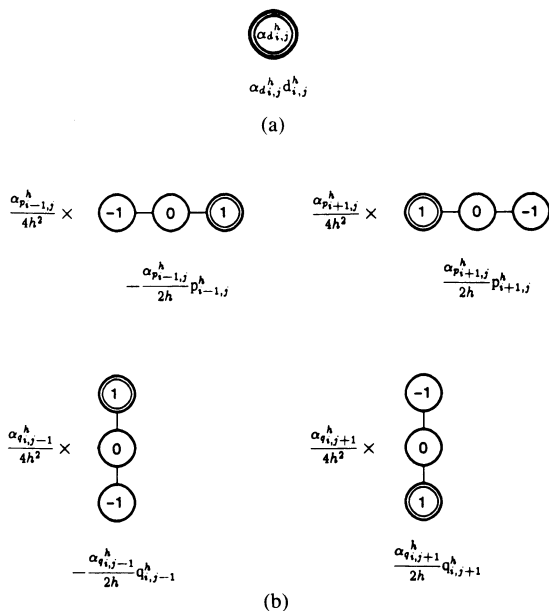


Fig. 5. (a) Depth constraint molecule. (b) Orientation constraint molecules.

the upper left molecule in Fig. 5(b) sums with the nodal molecule only if $(i - 1, j) \in P$, the upper right only if $(i + 1, j) \in P$, the lower left only if $(i, j - 1) \in Q$, and the lower right only if $(i, j + 1) \in Q$. Fig. 7 illustrates nodal molecules which can result from inhibition by discontinuities.

Discretizing $u(x, y)$, $\rho(x, y)$, and $\tau(x, y)$ over the same grid of nodes is convenient for expressing the nodal equation. However, it is more natural to position depth discontinuities on the links half way between nodes $\rho_{k,l}^h = \rho(kh + (h/2), lh + (h/2))$, since $u_{i,j}^h$ is undefined at a depth discontinuity. Orientation discontinuities can remain coincident with nodes $\tau_{i,j}^h = \tau(ih, jh)$ (since $u_{i,j}^h$ is defined at an orientation discontinuity). Generally, a discontinuity inhibits a molecule if it coincides with a constituent atom or link.

Variants of the above summation and inhibition rules




Fig. 6. Interior nodal molecules. (a) Molecule at node (i, j) away from constraints and discontinuities (here $\rho_{i,j}^h = \tau_{i,j}^h = 1$ and nodal equation is $(20/h^2)u_{i,j}^h - (8/h^2)(u_{i-1,j}^h + u_{i+1,j}^h + u_{i,j-1}^h + u_{i,j+1}^h) + (2/h^2)(u_{i-1,j-1}^h + u_{i+1,j-1}^h + u_{i-1,j+1}^h + u_{i+1,j+1}^h) + (1/h^2)(u_{i-2,j}^h + u_{i+2,j}^h + u_{i,j-2}^h + u_{i,j+2}^h) = 0$). (b) Molecule at interior orientation discontinuity node (i, j) (here $\tau_{i,j}^h = 0$ and nodal equation is $4u_{i,j}^h - u_{i-1,j}^h - u_{i+1,j}^h - u_{i,j-1}^h - u_{i,j+1}^h = 0$). Note that molecules represent h^2 times 13-point finite difference approximation to biharmonic (thin plate) operator and 5-point approximation to Laplacian (membrane) operator [1, p. 885].

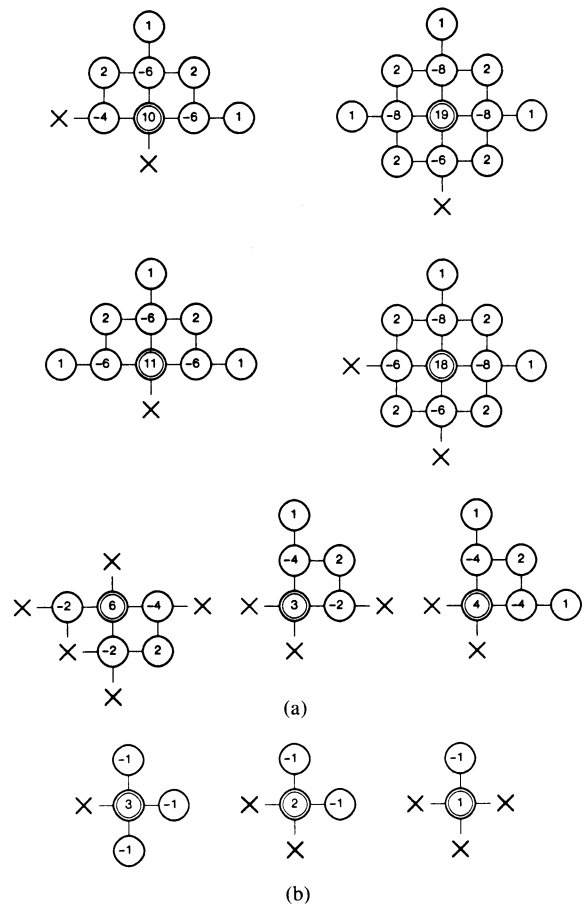


Fig. 7. Molecular inhibition at discontinuities. (a) Nodal molecules at boundary nodes (double circles) near depth discontinuity nodes (X's). (b) Nodal molecules at boundary orientation discontinuities (double circles) next to depth discontinuities (X's).

are possible. One useful alternative is to create horizontal or vertical depth discontinuities by introducing adjacent orientation discontinuities that inhibit only the “ $(-2)-(-2)$ ” molecules in Fig. 4.

C. Nodal Computations and Multiresolution Relaxation

The system of nodal equations features computationally desirable properties: its matrix becomes positive definite

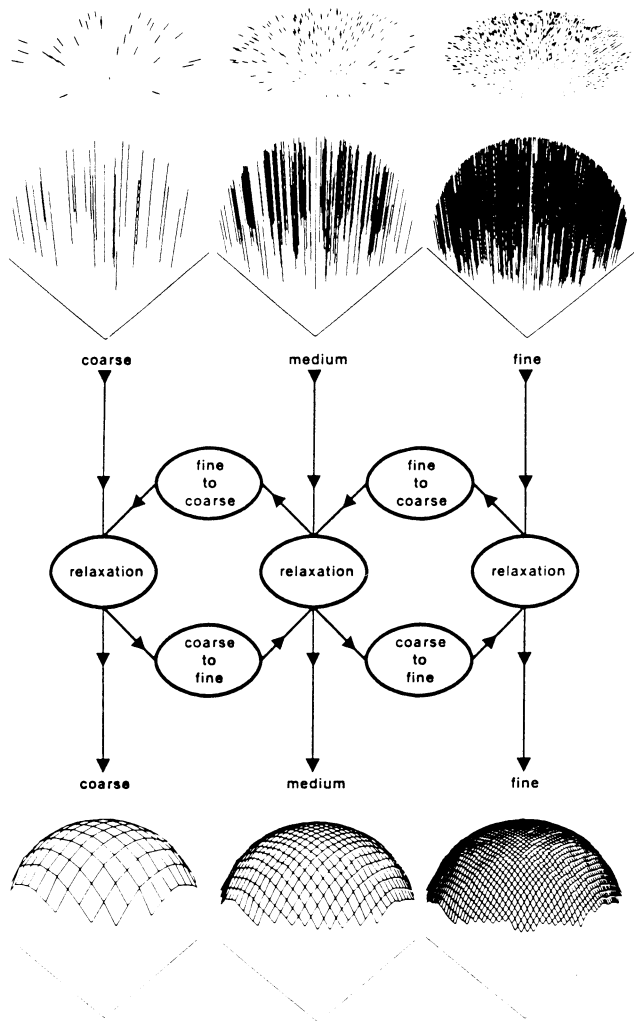


Fig. 8. Structure of the multiresolution surface reconstruction algorithm. Iterative relaxation processes propagate information within each level. Coarse-to-fine (prolongation) and fine-to-coarse (restriction) processes transfer information between levels. Synthetic orientation and depth constraints consistent with a hemispherical surface are input (top). Algorithm computes dense multiscale surface representation (bottom).

[for fixed $\rho(x, y)$ and $\tau(x, y)$] as soon as the available constraints satisfy the conditions for a well-posed problem. Moreover, the matrix is sparse (entries predominantly zero), banded, and symmetric, due to the local support of the finite element representation. However, its size $N^h \times N^h$ may become extremely large, since the number of pixels N^h in typical images can range from 10^4 to 10^6 or greater. This combination of properties suggests the application of iterative techniques which exploit sparsity, such as relaxation methods [32].

A parallel (Jacobi-type) *nodal relaxation computation* at node (i, j) may be written as follows:

$$u_{i,j}^{h(t+1)} = u_{i,j}^{h(t)} - \omega \left(\frac{\partial S_{\rho\tau}^h(\mathbf{u}^h)}{\partial u_{i,j}^h} + \frac{\partial \mathcal{G}^h(\mathbf{u}^h)}{\partial u_{i,j}^h} \right)^{(t)}, \quad (22)$$

where the bracketed superscripts contain the iteration index and ω is a (time) step size. The aforementioned molecular summation and inhibition rules fully automate the node-by-node construction of the bracketed term in this local-support computation. The computation is spatially

noninvariant and it changes temporally according to perturbations in the local structure of constraints and discontinuities. Gauss-Seidel nodal relaxation computations can be constructed similarly, as was done in [51].

For large problems the nodal relaxation computations usually suffer from very slow convergence. There exist efficient algorithms based on multigrid relaxation methods [31] for a number of visual problems [55], including interpolation problems involving globally continuous thin-plate splines [51]. The multiresolution relaxation approach is well suited to the visible-surface reconstruction process. Fig. 8 illustrates the structure of a three-level instance of the multiresolution surface reconstruction algorithm. A considerable simplification of the computations results from a 2:1 resolution reduction between adjacent levels (members of the hierarchy of embedded finite element subspaces). The hierarchy of representations and component processes are coordinated to increase computational efficiency. A recursive coordination strategy (see, [51], [55]) was employed in the experiments described below.

V. EXPERIMENTS WITH THE ALGORITHM

The multiresolution visible-surface reconstruction algorithm was tested on a variety of data sets including synthetic data, structured light (laser) rangefinder data, automated stereopsis and photometric stereo data from natural images, and digital terrain model data. This section presents selected results ([53] contains more examples and further details). For all the examples presented, the intralevel process was Gauss-Seidel relaxation and the algorithm was started from zero initial approximations on all the levels. Discontinuity maps have preset borders of depth discontinuities to introduce natural (free) boundary conditions.

A. Experiments Involving Preset Discontinuities

Synthetic Data: The first two examples illustrate multiresolution, piecewise continuous surface reconstructions from randomly placed depth constraints and prescribed discontinuities. Fig. 9 presents the reconstruction of a circular shell and Fig. 10 presents the reconstruction of stacked circular planes. Fig. 11 shows the reconstruction from orientation constraints of a pyramidal surface with orientation discontinuities. Fig. 12 presents the reconstruction of a hemispherical surface from scattered orientation constraints. Since it is impossible to determine absolute depth solely from orientation constraints, a relative-depth reconstruction results, with the center of gravity of the reconstructed surface resting near the $x - y$ plane. However, the integration of depth constraints with orientation constraints determines the absolute depth of the surface, as illustrated by Fig. 13. The constraint values in this latter example are corrupted by uniformly distributed noise. With the constraint parameters chosen, the surface is slightly bumpy on the finest level; the bumpiness can be reduced by decreasing the constraint parameter values (loosening the springs of the physical model).




Fig. 9. Multiresolution reconstruction of a shell from depth constraints. (a) Constraints obtained by randomly sampling at three resolutions a hemisphere with z values scaled by a radial sinusoid. Nodes outside circular region occupied by constraints were preset depth discontinuities. (b) Reconstructed surface representation. (Grid dimensions: $N_x^{h_1} = N_y^{h_1} = 17$, $N_x^{h_2} = N_y^{h_2} = 33$, $N_x^{h_3} = N_y^{h_3} = 65$. Grid spacings: $h_1 = 0.4$, $h_2 = 0.2$, $h_3 = 0.1$. Constraint density: 15 percent. Constraint parameters: $\alpha_d^{h_j} = 2.0/h_j$. Computation: 24.25 work units. A work unit is the amount of computation required per relaxation iteration on the finest level.)




Fig. 10. Multiresolution reconstruction of stacked circular planes from depth constraints. (a) Constraints obtained by randomly sampling the planar surfaces at three resolutions. Depth discontinuities are placed along circular arcs bounding the planes and along the outer edges of the grids. (b) Reconstructed surface representation. (Grid dimensions: $N_x^{h_1} \times N_y^{h_1} = 22 \times 17$, $N_x^{h_2} \times N_y^{h_2} = 43 \times 33$, $N_x^{h_3} \times N_y^{h_3} = 85 \times 65$. Grid spacings: $h_1 = 0.4$, $h_2 = 0.2$, $h_3 = 0.1$. Constraint density: 15 percent. Constraint parameters: $\alpha_d^{h_j} = 2.0/h_j$. Computation: 20.375 work units.)




Fig. 11. Multiresolution reconstruction of a pyramidal surface from orientation constraints. (a) Constraints at three resolutions—values are constant within each quadrant. Nodes along quadrant boundaries are preset orientation discontinuities. Outer nodes are preset depth discontinuities. (b) Reconstructed surface representation. (Grid dimensions: $N_x^{h_1} = N_y^{h_1} = 17$, $N_x^{h_2} = N_y^{h_2} = 33$, $N_x^{h_3} = N_y^{h_3} = 65$. Grid spacings: $h_1 = 0.4$, $h_2 = 0.2$, $h_3 = 0.1$. Constraint density: 100 percent. Constraint parameters: $\alpha_p^{h_j} = \alpha_q^{h_j} = 4.0/h_j$. Computation: 19.5 work units.)




Fig. 12. Multiresolution of a hemisphere from orientation constraints. (a) Constraints randomly sampled at three resolutions from a hemisphere. Nodes outside the hemispherical surface patch were preset depth discontinuities. (b) Reconstructed surface representation. (Grid dimensions: $N_x^{h_1} = N_y^{h_1} = 17$, $N_x^{h_2} = N_y^{h_2} = 33$, $N_x^{h_3} = N_y^{h_3} = 65$. Grid spacings: $h_1 = 0.4$, $h_2 = 0.2$, $h_3 = 0.1$. Constraint density: 30 percent. Constraint parameters: $\alpha_p^{h_j} = \alpha_q^{h_j} = 4.0/h_j$. Computation: 22.125 work units.)




Fig. 13. Multiresolution reconstruction of the hemisphere from depth and orientation constraints. (a) Depth constraints and (b) orientation constraints consistent with a hemisphere at three resolutions. (c) Reconstructed surface representation; note the absolute height above the base plane is obtained. (Grid dimensions: $N_x^{h_1} = N_y^{h_1} = 17$, $N_x^{h_2} = N_y^{h_2} = 33$, $N_x^{h_3} = N_y^{h_3} = 65$. Grid spacings: $h_1 = 0.4$, $h_2 = 0.2$, $h_3 = 0.1$. Constraint density: 15 percent. Constraint noise: 10 percent uniform. Constraint parameters $\alpha_d^{h_j} = 2.0/h_j$, $\alpha_p^{h_j} = \alpha_q^{h_j} = 4.0/h_j$. Computation: 17.75 work units.)

Structured Light Data: The multiresolution algorithm was applied to the reconstruction of several objects from raw range data supplied by a laser rangefinder constructed at MIT by P. Brou. The scan resolution in the y direction is half that in the x direction. The examples involve a four-level surface reconstruction algorithm. The raw rangefinder data were introduced as depth constraints at the finest level and transferred to the coarser levels by successive 2×2 averaging between levels. To expediently segment the objects from the supporting platform, values smaller than a threshold were treated as depth discontinuities. Fig. 14 shows the reconstructed surface of a light-bulb. The algorithm smooths the noise in the data and reconstructs the missing points.

Photometric Stereo Data: The surface reconstruction algorithm provides a noise resistant technique for computing depth from the surface orientation data provided by photometric stereo [62]. This is demonstrated using the image of a toroid (Fig. 15). The photometric stereo data shown (generated by a program implemented at MIT by K. Ikeuchi) were introduced as orientation constraints on a two-level algorithm. Aside from sporadic missing data, the constraints on the coarse level are dense, while only every other node on the fine level is a constraint. Fig. 15(c) shows the reconstructed toroid.

Correlation-Based Stereo Data: Fig. 16(a) shows a stereopair that was input to the correlation-based stereo program described in [36]. Fig. 16(b) shows the output of the program. The neutral gray patches indicate regions of unknown disparity, where the algorithm has failed to produce a match. Successively averaging by factors of two, subsampled versions of the disparity data on the finest






Fig. 14. Reconstruction of a light bulb from range data. (Finest grid dimensions: $N_x^{h_4} \times N_y^{h_4} = 257 \times 281$. Grid spacings: $h_1 = 0.8$, $h_2 = 0.4$, $h_3 = 0.2$, and $h_4 = 0.1$. Constraint parameters: $\alpha_d^{h_j} = 0.2/h_j$. Computation: 9.78 work units.)



(a)



(b)

(c)

Fig. 15. Reconstruction of a torus from photometric stereo data. (a) Image of a matte white toroid. (b) Orientation constraints provided by photometric stereo. (c) Reconstructed torus. (Grid dimensions $N_x^{h_1} = N_y^{h_1} = 51$ and $N_x^{h_2} = N_y^{h_2} = 101$. Constraint parameters: $\alpha^{h_j} = 4.0/h_j$. Computation: 52.0 work units.)

level were introduced as input to the multiresolution algorithm on three coarser levels. Relatively small constraint parameter values were chosen in order to counteract the false matches and noise present in the disparity data. Fig. 16(c) shows the reconstructions on the three coarsest levels as 3-D plots (the finest level was too dense to render as a wire frame surface). Fig. 16(d) shows isoelevation contour maps of the solution on all levels.

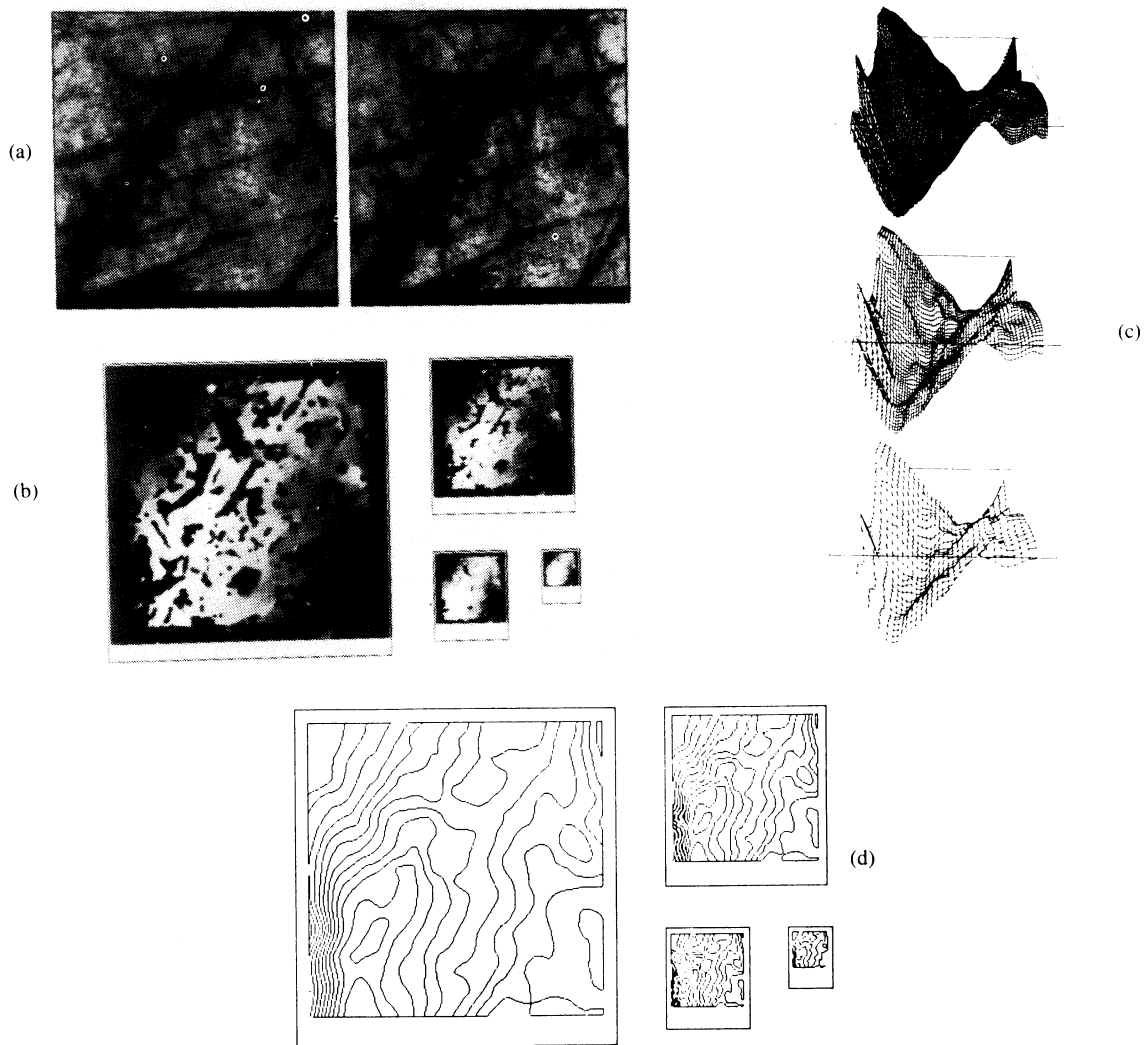


Fig. 16. Reconstruction of terrain from correlation-based stereo data. (a) Natural terrain stereopair (images: 256×256 pixels, quantized to 256 gray levels; courtesy of U.S. Defense Mapping Agency). (b) Output of correlation-based stereo program (brightness proportional to disparity; disparity unknown in neutral grey regions). (c) Reconstructed terrain on three coarsest levels. (d) Isoelevation contour maps. (Grid dimensions: $N_x^{h_1} = N_y^{h_1} = 33$, $N_x^{h_2} = N_y^{h_2} = 65$, $N_x^{h_3} = N_y^{h_3} = 129$, $N_x^{h_4} = N_y^{h_4} = 257$. Grid spacings: $h_1 = 0.8$, $h_2 = 0.4$, $h_3 = 0.2$, $h_4 = 0.1$. Constraint parameters: $\alpha_d^{h_i} = 0.01/h_i^2$. Computation: 29.0 work units.)

Feature-Based Stereo Data: The stereopair of Fig. 17(a) was input to a three-channel version of the (zero-crossing) feature-based stereo program described in [29]. Fig. 17(b) shows the output of the program. Disparity information is available only along zero crossing contours at the three finest scales. This disparity data were input to a four-level surface reconstruction algorithm, with the constraints on the coarsest level derived by averaging the constraints from the next finer level. Fig. 17(c) shows the reconstructions on the three coarsest levels as 3-D plots. Fig. 17(d) shows isoelevation contour maps of the solution on all levels.

Digital Terrain Map Data: Fig. 18 presents the application of a four-level surface reconstruction algorithm to contoured terrain elevation data. The data was generated by J. Mahoney, using a digitizing tablet to trace manually

the isoelevation contours in a map of the Black River Gorges (published by the UK Ministry of Defense). Fig. 18(b) shows contour plots of the reconstructed terrain on all levels. The elevations of the reconstructed contours are chosen halfway between the original constraint contours of Fig. 18(a) to avoid favorable bias. The reconstructed contours are somewhat smoother than the contours in the published map. The jaggedness introduced by manual digitization has been reduced; the constraint parameters regulate the smoothing. Fig. 18(c) shows shaded images of the reconstructed terrain. Comparison of terrain reconstructions using the controlled-continuity model against reconstructions using a simpler membrane spline model (Laplacian smoothing) indicates that the latter generally suffers from insufficient smoothness and produces flat spots across terrain peaks (see also [9]).

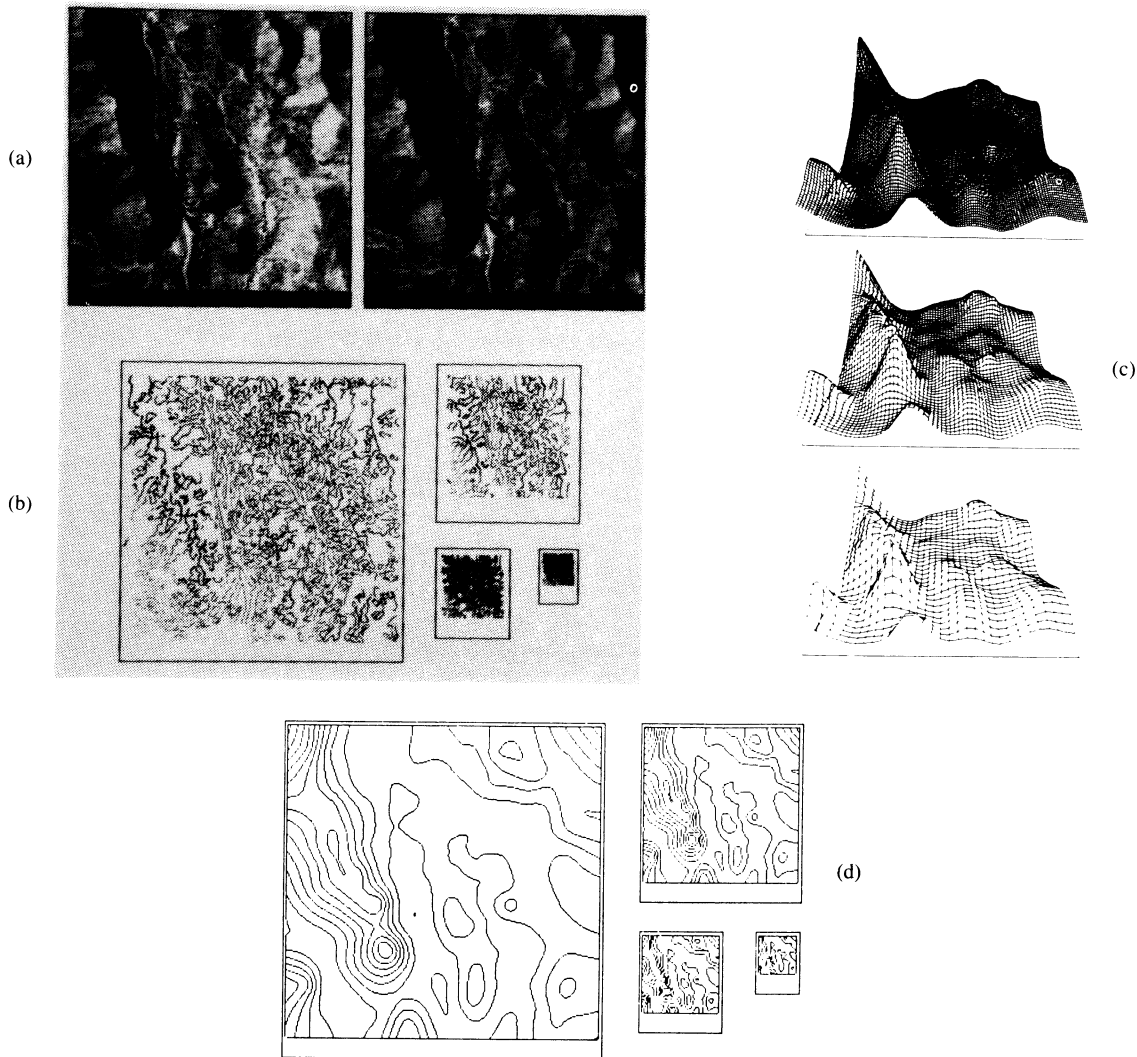


Fig. 17. Reconstruction of terrain from feature-based stereo data. (a) Natural terrain stereopair (images: 512×512 pixels, quantized to 256 levels; courtesy of US Army Engineer Topographic Labs). (b) Output of feature-based stereo program (contour intensity proportional to disparity). (c) Reconstructed terrain on three coarsest levels. (d) Isoelevation contour maps. (Grid dimensions: $N_x^{h_1} = N_y^{h_1} = 33$, $N_x^{h_2} = N_y^{h_2} = 65$, $N_x^{h_3} = N_y^{h_3} = 129$, $N_x^{h_4} = N_y^{h_4} = 257$. Grid spacings: $h_1 = 0.8$, $h_2 = 0.4$, $h_3 = 0.2$, $h_4 = 0.1$. Constraint parameters: $\alpha_d^{h_j} = 0.01/h_j^2$. Computation: 31.0 work units.)

B. Discontinuity Identification Experiments

When all discontinuities are preset, the discontinuity map, $\rho_{i,j}^h = \rho(ih, jh)$ and $\tau_{i,j}^h = \tau(ih, jh)$, comprises a fixed part of the input data. More generally, some or all of the discontinuities are unknown, so the discontinuity map undergoes refinement during surface reconstruction. The following experiments involve the automatic identification of surface discontinuities, first by local validation, followed by variational continuity control (refer to Section III).

Local Validation: Fig. 19 illustrates the application of the local validation method to a random dot stereogram. Fig. 19(b) shows the depth constraints generated by a three-channel version of the feature-based stereo program [29]. The discontinuities identified using one cycle of lo-

cal validation are superimposed onto the final disparity maps in Fig. 19(c). Fig. 19(d)–(g) illustrates the steps of the local validation method in more detail. Fig. 20 illustrates a similar experiment using an aerial view stereopair. The results demonstrate the feasibility of identifying significant discontinuities using local validation, but the interpolation process can be seen “leaking” through gaps due to undetected discontinuities. No single value of the global limit t_d can be expected to produce perfect results, especially with natural imagery. As was suggested in Section III-A, however, decreasing the disparity gradient limit in a sequence of cycles, iterating to near-equilibrium each time, will elicit increasingly shallow discontinuities.

Variational Continuity Control: Variational continuity control involves a similar multistep strategy guided by the energy functional $\mathcal{E}(v, \rho, \tau)$ in VP2. The present imple-




Fig. 18. Reconstruction of digital terrain map data. (a) 256×256 digital contour array input to algorithm (contour brightness proportional to elevation; local averaging of finest grid yielded constraints on coarser grids; patch to lower right indicates a lake). (b) Reconstructed isoelevation contour map. (c) Shaded terrain representations. (Grid Dimensions: $N_x^{h_1} = N_y^{h_1} = 33$, $N_x^{h_2} = N_y^{h_2} = 65$, $N_x^{h_3} = N_y^{h_3} = 129$, $N_x^{h_4} = N_y^{h_4} = 257$. Grid spacings: $h_1 = 0.8$, $h_2 = 0.4$, $h_3 = 0.2$, $h_4 = 0.1$. Constraint parameters: $\alpha_d^{h_j} = 0.5/h_j^2$.)

mentation employs the following discrete form for the discontinuity functional $\mathcal{D}(\rho, \tau)$ in (10):

$$\mathcal{D}^h(\rho^h, \tau^h) = \sum_{i,j} [\beta_d^h D_{i,j}^h(\rho^h) + \beta_o^h O_{i,j}^h(\tau^h)], \quad (23)$$

where $D_{i,j}^h$ and $O_{i,j}^h$ are relative potential energies of local discontinuity configurations in the depth and orientation discontinuity maps, respectively. Fig. 21 illustrates some configurations and associated energies chosen heuristically to favor the formation of continuous and smoothly curving discontinuity contours; higher energies are assigned to isolated discontinuities, terminations, sharp bends, junctions, and discontinuity clumps [25].

The strategy for "solving" VP2 is as follows: The outer iteration consists of 1) each continuity control parameter $\rho_{i,j}^h$ or $\tau_{i,j}^h$ (in parallel) flips its value in $\{0, 1\}$ if this reduces the energy \mathcal{E}^h (10), then 2) given the updated $\rho_{i,j}^h$ or

$\tau_{i,j}^h$, the basic surface reconstruction algorithm obtains the unique minimum of the inner functional (11). First, the outer iteration identifies depth discontinuities, with β_d^h initially set to a high value (to heavily penalize fracture), then gradually lowered. Next, the outer iteration identifies orientation discontinuities with β_o^h decreasing in the same way. Orientation discontinuity identification is delayed, otherwise the small surface inflections occurring near undetected depth discontinuities (see Fig. 3) are easily misconstrued as orientation discontinuities, which retards the optimization process. Fig. 22 illustrates the variational continuity control approach using a synthetic example.

The above describes a continuation procedure [61] for solving VP2, in which β_d^h and β_o^h serve as continuation variables. VP2 becomes increasingly convex—i.e., simpler—for larger values of these variables. By gradually

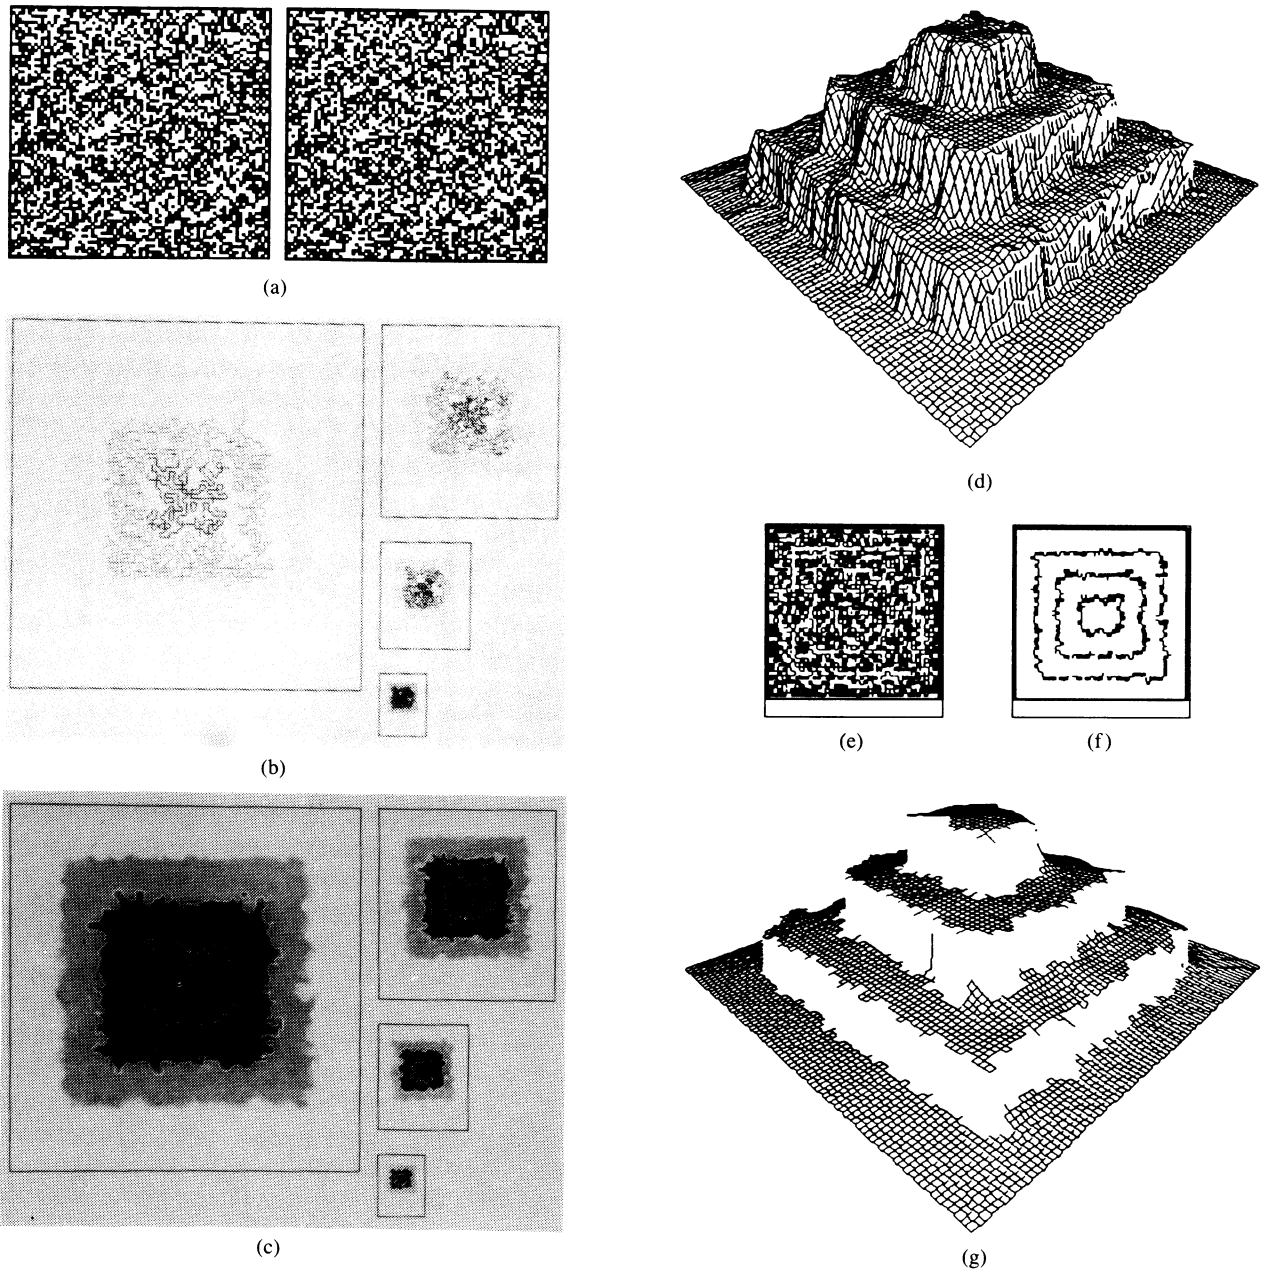


Fig. 19. Reconstruction of a random dot stereogram by local validation. (a) Synthesized stereogram of four planar surfaces stacked in depth. (b) Depth constraints for random dot stereogram (finest level dimensions: 320×320 ; constraints on coarsest level obtained by averaging from finer level). (c) Piecewise continuous disparity maps with detected discontinuities (white contours) superimposed. (d)-(g) Local validation method on a coarse level. (d) Reconstructed C^1 surface. (e) Zero crossings of discrete bending moment $M_{i,j}^h = -1/h^2(u_{i-1,j}^h + u_{i+1,j}^h + u_{i,j-1}^h + u_{i,j+1}^h - 4u_{i,j}^h)$ (black points). (f) Significant zero crossings for which $G_{i,j}^h > t_d$ with $t_d = 1$, where $G_{i,j}^h = 1/4h^2[(u_{i+1,j}^h - u_{i-1,j}^h)^2 + (u_{i,j+1}^h - u_{i,j-1}^h)^2]$. (g) Piecewise continuous surface results from inserting latter into discontinuity map (by setting associated $\rho_{i,j}^h$ to zero) and continuing iterative reconstruction process.

decreasing β_d^h and β_o^h , a trajectory is followed from the optimum of the simple problem towards a satisfactory "solution" to the difficult, nonconvex target problem (β_d^h and β_o^h small). In general, relatively few of the $\rho_{i,j}^h$ or $\tau_{i,j}^h$ are observed to change at each step, so few iterations are necessary per step (in a uniprocessor implementation,

computation may be concentrated at nodes near modified discontinuities). Increasingly accurate discontinuities emerge as β_d^h and β_o^h decrease. Once a true discontinuity appears it tends to persist (a hysteresis effect); thus, beyond some minimal β_d^h and β_o^h values, a slight increase will eliminate some spurious discontinuities and release




Fig. 20. Local validation applied to piecewise continuous stereo reconstruction. (a) Aerial view stereopair of a hospital complex (images courtesy of U. British Columbia, Faculty of Forestry). (b) Disparity contour on finest level (size 320×320) generated by feature based stereo program (darkness proportional to disparity). (c) Full disparity map generated by surface reconstruction algorithm on finest level. (d) Discontinuities identified on finest level (white contours) superimposed on reconstructed piecewise continuous disparity map.

additional energy [this occurs in the transition from Fig. 22(f) to (g)]. Although it has achieved the global optimum of $VP2$ in Fig. 22(h), the continuation procedure can generally be expected to yield good, although not necessarily optimal solutions. Its deterministic nature and efficiency are attractive.

VI. DISCUSSION

This paper has developed a computational theory of visible-surface representations. The visible-surface reconstruction process which implements the theory provides a uniform treatment of integration, interpolation, discontinuities, and efficiency, the four computational goals outlined in the introduction. The concluding discussion surveys research directly relevant to the computation of




Fig. 21. Some local discontinuity configurations and associated relative energies. Circles represent nodes (i, j) , while X's denote discontinuities (positions where ρ^h or τ^h are 0; depth discontinuities occur on links between nodes while orientation discontinuities coincide with nodes.) (a) Energies $D_{i,j}^h$ for depth discontinuity configurations (cf. [25]). (b) Energies $O_{i,j}^h$ for some orientation discontinuity configurations. Rotating configurations by increments of 90 degrees leaves associated energies unaffected.

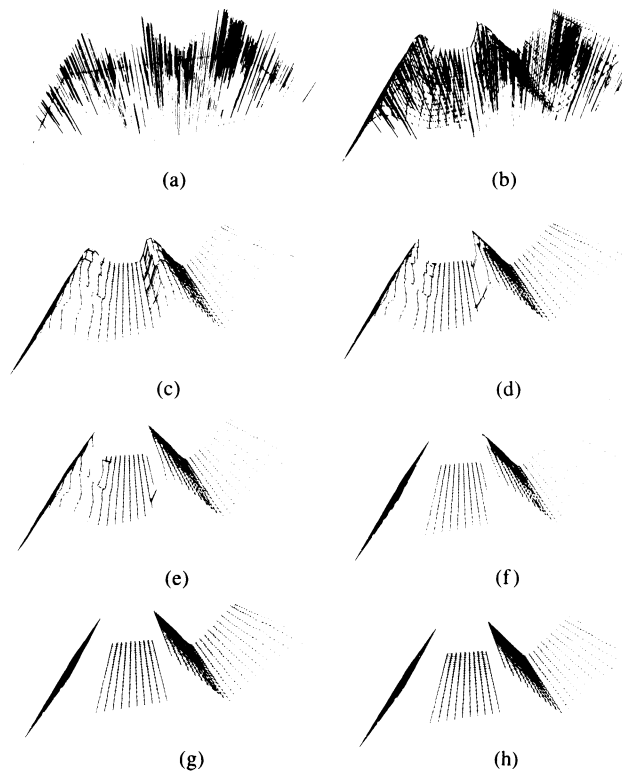


Fig. 22. Variational continuity control method for piecewise continuous reconstruction. (a) Scattered depth constraints consistent with sloping planes meeting discontinuously. (b) Globally C^1 reconstructed surface obtained with high β_d^h and β_o^h . Note smeared depth discontinuities and rounded orientation discontinuities. (c)-(h) Evolution of the variational continuity control process. β_d^h is lowered first to identify depth discontinuities (c)-(g) β_o^h is lowered next to identify orientation discontinuities. (h) Final piecewise continuous surface.

visible-surface representations and proposes future research directions.

A. Survey

Barrow and Tenebaum [3], [4] describe an approach to reconstructing smooth surfaces from noisy visual data. Their algorithms apply only to a restricted class of surfaces, but the intrinsic image model [2] and many of their ideas regarding its computation have influenced subsequent work significantly. Grimson [27], [28] employs an established surface fitting method for the smooth interpolation of visual surfaces from depth constraints. The method involves the minimization of a particular functional containing second derivatives (Grimson referred to it as the “quadratic variation,” but it is better known in the literature as the thin-plate spline) [47]. Brady and Horn [12] point out that the resulting biharmonic interpolant characterizes the bending energy of a thin plate (see also [47]). The thin-plate reconstruction model is developed further by Terzopoulos [51]. The high quality of thin-plate interpolating surfaces relative to competing techniques are responsible for their popularity in engineering applications (see [47], [56] for references).

A serious drawback of the thin-plate spline, however, is that its global C^1 continuity gives rise to undesirable overshoots near large gradients. Terzopoulos [53] proposes to ameliorate the large gradient problem by introducing global tension into the surface (controlled by a single tension parameter), in analogy to cubic splines in tension (see also the independent work of Franke [23]). Another potential problem is the viewpoint noninvariance of the linear thin-plate approximation, which becomes significant when large gradients are imposed by very sparse data [7]. One possible countermeasure is to invoke the large deflection theory of thin plates [53], [7]. Blake and Zisserman [8] have made progress against the computational difficulties inherent to this nonlinear model. Alternatively, viewpoint invariance results without sacrificing linearity when the surface is represented parametrically as $[x(u, v), y(u, v), z(u, v)]$; here, the intrinsic variables u and v no longer correspond to the image coordinates x and y .

Problems with large gradients are best eliminated by an explicit treatment of surface discontinuities. The introduction of depth (jump) discontinuities has been attempted with regard to harmonic and biharmonic interpolation of geophysical data [9] (see also [14]) and with regard to distance weighted local surface approximation for CAGD [24], but irregularly shaped discontinuities pose difficulties for these schemes. Terzopoulos [52] presents an approach to piecewise continuous surface reconstruction from both depth and orientation constraints and (possibly irregular) discontinuities which is based on composite variational principles involving local thin-plate and membrane elements.

These composite variational principles generalize to the class of controlled-continuity spline models [56], of which $S_{\rho\tau}$ in equation (5) is an instance. Controlled-continuity

spline models include as special cases the piecewise smoothness constraints employed by Blake [6] and Geman and Geman [25] for image restoration (dense data) and by Marroquin [42] and Koch *et al.* [37] for surface reconstruction (the models in [25], [42] are cast in terms of Markov random fields). In particular, the functionals in [6], [25] correspond to piecewise zeroth-order (constant) splines, while those in [37], [42] correspond to piecewise first-order (harmonic or membrane) splines. However, a second-order surface model, such as $S_{\rho\tau}$ is the lowest order capable of integrating depth and orientation constraints and discontinuities.

The surface reconstruction examples presented by Terzopoulos in [52], [53] employ the vocal validation method for discontinuity detection (Section III-A). Langridge [38] suggests a related scheme for the detection of surface orientation discontinuities based on curvature peaks. Grimson and Pavlidis [30] detect discontinuities prior to reconstruction by monitoring local statistics of the residuals of a local planar approximation to the depth data. Brady *et al.* [13] use 1-D scale space methods along lines of curvature of a reconstructed surface to find discontinuities.

Blake introduces in [6] the idea of including a discontinuity penalty term as part of the minimal energy formulation of the piecewise continuous reconstruction problem. Geman and Geman [25] introduce an explicit line process whose role is to favor continuous image intensity edge configurations. Similar local edge compatibility rankings favoring continuity have been employed in connection with relaxation labeling curve enhancement processes [64], [65]. Marroquin [42] adopts the line process idea for depth discontinuity detection in surface reconstruction. The present paper makes use of a line process in the variational continuity control approach (see Fig. 21). In view of the underlying physical model, the technique is related to certain computational methods in fracture mechanics for simulating the release of energy along fractures propagating in solids [63]. Rather than toggling binary discontinuity variables in a line process, Koch *et al.* [37] employ a continuous, nonlinear gain mechanism to insert depth discontinuities. This mechanism, due to Hopfield, can be adapted to continuously estimate $\rho(x, y)$ and $\tau(x, y)$ in VP2. Also of interest is the work of Mumford and Shah [43], who propose a method (implemented in 1-D) to solve the Euler-Lagrange equations for smooth boundaries.

The conventional numerical construction of thin-plate surfaces and other generalized splines is in terms of their reproducing kernel Hilbert space representation [47]. This requires the solution of dense systems of equations of size proportional to the number of point constraints (see [56, Appendix D] for a summary formulation). Boulton and Kender [10] pursue reproducing kernel techniques for surface reconstruction. Advantages accrue for globally continuous reconstruction when there are relatively few constraints or when each constraint is introduced sequentially. It is possible, in principle, to supplement the reproducing kernel spaces with step functions to treat

piecewise continuous reconstruction, although the formulation presents difficulties, especially with regard to estimating irregular discontinuities in multidimensions [48]. In view of the requirements specific to the computation of visible-surface representations—very large constraint sets, undetermined discontinuities, desirability of massively parallel implementation, etc.—it appears preferable to apply finite element techniques which yield sparse systems of equations, as is done in the present paper.

Barrow and Tenenbaum [3], [4] use relaxation while Grimson [27], [28] uses standard optimization algorithms for surface interpolation. These iterative algorithms have parallel variants which are thought to be “biologically feasible.” However, their tendency to be poorly conditioned, even for reconstruction problems of moderate size, results in very slow convergence. Excruciatingly slow convergence also limits the size of problem that can be attempted with optimization via simulated annealing [25], [42]. Terzopoulos [51] shows that multiresolution surface reconstruction based on multigrid relaxation methods accelerates convergence dramatically while maintaining biological feasibility. Cochran and Medioni [19] describe another implementation of this multiresolution surface reconstruction algorithm. An efficient multigrid algorithm for reconstructing surfaces from shaded monocular images is described in [55]. A concurrent multigrid coordination strategy which solves a coupled, multilevel version of the variational principle is developed in [54]; the approach is well suited to massively parallel hardware because it maintains simultaneous processor activity in all levels.

B. Research Directions

Alternative organizations of shape information in visible-surface representations are worthy of further study. One possible variant is a *relative* shape representation, where only the coarsest grid contains absolute depth and orientation values, while each finer grid contains increasingly detailed perturbations relative to the sum of these values over all coarser grids. Surface structure parsimoniously decomposes into the hierarchy of finite element subspaces of the multiresolution representation (roughly analogous to a Fourier decomposition). A relative shape representation may simplify the multigrid relaxation algorithm by enabling the levels to run virtually independently of one another (see [53, Ch. 11] for a detailed discussion). Another representation of interest is the explicit depth/slope representation proposed in [33].

The coupling of visual constraints to the surface model requires further analysis. The surface model can be related to expectations regarding the class of admissible surfaces, and a connection exists between the constraint parameters α_i and the statistical properties of noise or inaccuracy in the data [56]. Assuming that low-level visual processes can associate a confidence or variance estimate with each shape constraint, it appears possible to systematically assign appropriate values to the constraint param-

eters. In particular, cross validation techniques [60] may provide a means of setting λ^{-1} [refer to text after (7)] to optimally tune the smoothness of the reconstructed surface according to the noise present in the data.

The arbitration of visible-surface reconstruction by higher-level processes has received insufficient attention. Among other important functions, an arbitrator could deal with occasional outliers or massively inconsistent data by nullifying individual or entire sets of constraint parameters. The straightforward treatment of a particular kind of rivalry, that due to transparent surfaces is described in [53, Ch. 11]. The arbitrator monitors the approximation error between reconstructed surface and shape constraints over broad areas. An excessive error triggers a grouping process which clusters constraints. Multiple surfaces are then reconstructed over the same region, one for each homogeneous constraint population.

The line processes used in the discontinuity detection experiments are rather primitive. Zucker and Parent [65] employ a more sophisticated discrete encoding of local orientation in a relaxation labeling algorithm for contour detection in images. It appears possible to employ a similar encoding in the estimation of surface discontinuities by variational continuity control via *VP2*. A continuous encoding of the local orientations of curvilinear boundary elements appears even more desirable, but introduces additional real-valued nodal parameters in the discontinuity map.

Analog computation by electrical networks [35] is an attractive approach to generating visible-surface representations. The major advantages of large analog networks are fault tolerance, noise dissipation, and very rapid settling to steady-state (within a few network time constants). Lumped analog networks may be designed systematically whose steady-state voltages and currents represent quantities of interest in variational formulations of visual problems such as surface reconstruction [53], [44], [37], [33]. The network design process parallels finite element discretization; one or more electrical devices may be used to simulate the physical properties of each finite element [63]. Resistance networks for computing first order and second order spline interpolants are shown in Fig. 23.

It remains to explore ensuing processes that generate stable higher-level shape descriptions which are better tuned to object recognition. The visible-surface representation comprises only an intermediate, viewer-centered description of the 3-D surfaces in scenes, and the goal of subsequent processing is to abstract a rich set of object-centered features that are stable through viewpoint changes. This goal begins with *visible-surface analysis*, which is facilitated by the dense, quantitative shape information provided by visible-surface representations. This topic is considered further.

A promising approach to visible-surface analysis is to apply concepts from differential geometry [20]. For instance, according to the fundamental theorem of the local theory of surfaces (Bonnet), the analytic study of a sur-




Fig. 23. Analog networks for harmonic and biharmonic interpolation. (a) Network of resistances $r = 1/h^2$ solves Poisson's equation for displacement (depth) over a uniform grid of spacing h . Currents $g_{i,j}$ determined by the right-hand side of the equation are injected into each node. Voltages are applied to certain nodes in accordance with the boundary conditions and constraints of the problem. Discontinuities may be inserted by "breaking" resistors. (b) Two cascaded networks solve the biharmonic equation over a uniform grid. Top network solves a Poisson equation for the bending moment, while bottom network solves a similar equation for the displacement.

face consists of the study of its two fundamental forms; i.e., the six fundamental tensor coefficients (not all independent) as functions of the two independent param-




Fig. 24. Computing intrinsic and extrinsic surface properties in the visible-surface representation. (a) Computed Gaussian curvature $K_{i,j}$ of reconstructed lightbulb at four scales. Elliptic points ($K > 0$) are white, hyperbolic points ($K < 0$) are black, and parabolic points ($K = 0$) separate regions. (b) Computed principal direction field for reconstructed lightbulb at the two coarsest scales showing directions of greatest curvature (left) and least curvature (right).

ters of the surface. The fundamental forms are invariant under changes in surface parameterization, and together they determine surface shape up to rigid body transformations, making them ideal foundations for object-centered representations of surfaces [53, Ch. 11].

The finite element shape representation reduces the computation of fundamental forms and derived surface features such as the Gaussian, mean, and principal curvatures to the evaluation of simple algebraic expressions of neighboring nodal variables [see Appendix, (25)–(28)]. It therefore becomes straightforward to compute the elliptic, hyperbolic, parabolic, umbilic, and planar points, as well as geodesics, asymptotes, and lines of curvature. Fig. 24 illustrates some of these computations using the reconstructed surface of the lightbulb. The results demonstrate the feasibility of reliably computing from visible-surface representations higher-order intrinsic and extrinsic properties of surface shape (see also, e.g., [12], [22], [5], [59]). The reliability is attributable to the controlled-continuity surface model which overcomes the potentially

detrimental effects of noise in the data without destroying discontinuities.

APPENDIX

Denote a vector in \mathbb{R}^3 by $\mathbf{x} = [x^1, x^2, x^3]$, and a surface patch by $\mathbf{x} = \mathbf{x}(u^1, u^2) = [x^1(u^1, u^2), x^2(u^1, u^2), x^3(u^1, u^2)]$, where (u^1, u^2) is a point in parameter space. Let $\mathbf{x}_1 = \partial \mathbf{x} / \partial u^1$, $\mathbf{x}_2 = \partial \mathbf{x} / \partial u^2$, $\mathbf{x}_{12} = \partial^2 \mathbf{x} / \partial u^1 \partial u^2$, etc. A tangent vector is given by $d\mathbf{x} = \mathbf{x}_1 du^1 + \mathbf{x}_2 du^2 = \mathbf{x}_i du^i$, where the Einstein summation convention is used for an index occurring both as a superscript and subscript in a product. The first fundamental form is $I = d\mathbf{x} \cdot \mathbf{x} = \mathbf{x}_i \cdot \mathbf{x}_j du^i du^j = g_{ij} du^i du^j$, where $g_{ij} = \mathbf{x}_i \cdot \mathbf{x}_j$ for $i, j = 1, 2$ are the first fundamental (metric) tensor coefficients. The unit normal vector is $\mathbf{n} = \mathbf{x}_1 \times \mathbf{x}_2 / |\mathbf{x}_1 \times \mathbf{x}_2|$, and its differential is the vector $d\mathbf{n} = \mathbf{n}_i du^i$. The second fundamental form is $II = d^2 \mathbf{x} \cdot \mathbf{n} = \mathbf{x}_{ij} \cdot \mathbf{n} du^i du^j = b_{ij} du^i du^j$, where $b_{ij} = \mathbf{x}_{ij} \cdot \mathbf{n}$ for $i, j = 1, 2$ are the second fundamental tensor coefficients.

For the viewer-centered parameterization $u^1 = x^1 = x$ and $u^2 = x^2 = y$, the quadratic finite element $v(x, y) |_E$ is expressed as

$$\mathbf{x}(x, y) = [x, y, ax^2 + by^2 + cxy + dx + ey + f], \quad (24)$$

where the coefficients a to f are uniquely determined by (15) in terms of six unisolvent nodal variables associated with each element [51]. Hence,

$$\begin{aligned} \mathbf{x}_1 &= [1, 0, 2ax + cy + d]; \\ \mathbf{x}_2 &= [0, 1, 2by + cx + e]; \quad \mathbf{x}_{11} = [0, 0, 2a]; \\ \mathbf{x}_{22} &= [0, 0, 2b]; \quad \mathbf{x}_{12} = \mathbf{x}_{21} = [0, 0, c]; \end{aligned}$$

$$\begin{aligned} \mathbf{n} &= \frac{\mathbf{x}_1 \times \mathbf{x}_2}{|\mathbf{x}_1 \times \mathbf{x}_2|} \\ &= \frac{-[2ax + cy + d, 2by + cx + e, -1]}{\sqrt{(2ax + cy + d)^2 + (2by + cx + e)^2 + 1}}; \end{aligned}$$

yielding the first fundamental tensor coefficients

$$\begin{aligned} g_{11} &= \mathbf{x}_1 \cdot \mathbf{x}_1 = 1 + (2ax + cy + d)^2; \\ g_{22} &= \mathbf{x}_2 \cdot \mathbf{x}_2 = 1 + (2by + cx + e)^2; \\ g_{12} &= g_{21} = \mathbf{x}_1 \cdot \mathbf{x}_2 = (2ax + cy + d) \\ &\quad \cdot (2by + cx + e); \end{aligned}$$

and the second fundamental tensor coefficients

$$\begin{aligned} b_{11} &= \mathbf{x}_{11} \cdot \mathbf{n} = \frac{2a}{|\mathbf{x}_1 \times \mathbf{x}_2|}; \\ b_{22} &= \mathbf{x}_{22} \cdot \mathbf{n} = \frac{2b}{|\mathbf{x}_1 \times \mathbf{x}_2|}; \\ b_{12} &= b_{21} = \mathbf{x}_{12} \cdot \mathbf{n} = \frac{c}{|\mathbf{x}_1 \times \mathbf{x}_2|}, \end{aligned}$$

where

$$|\mathbf{x}_1 \times \mathbf{x}_2| = \sqrt{(2ax + cy + d)^2 + (2by + cx + e)^2 + 1}.$$

The fundamental tensor coefficients are therefore simple algebraic expressions over the element domain; however, at the central node $(0, 0)$ of the element their values reduce to

$$\begin{aligned} g_{ij}(0, 0) &= \begin{pmatrix} 1 + d^2 & de \\ de & 1 + e^2 \end{pmatrix}, \\ \text{and } b_{ij}(0, 0) &= (1 + d^2 + e^2)^{-1/2} \\ &\quad \cdot \begin{pmatrix} 2a & c \\ c & 2b \end{pmatrix}. \end{aligned}$$

The normal curvature in the direction $d\mathbf{u} = [du^i, du^j]^T$ is $\kappa_n = II/I$. Let κ_1 and κ_2 be the principal curvatures. The Gaussian curvature is

$$K = \kappa_1 \kappa_2 = \frac{b_{11} b_{22} - b_{12}^2}{g_{11} g_{22} - g_{12}^2} = \frac{4ab - c^2}{(1 + d^2 + e^2)^2}, \quad (25)$$

and the mean curvature is

$$\begin{aligned} H &= \frac{\kappa_1 + \kappa_2}{2} = \frac{b_{11} g_{22} - 2b_{12} g_{12} + b_{22} g_{11}}{g_{11} g_{22} - g_{12}^2} \\ &= \frac{2a(1 + e^2) - 2cde + 2b(1 + d^2)}{(1 + d^2 + e^2)^{3/2}} \end{aligned} \quad (26)$$

at the central node. Applying the differential equation for the lines of curvature,

$$\begin{aligned} (g_{11} b_{12} - g_{12} b_{11})(du^1)^2 + (g_{11} b_{22} - g_{22} b_{11}) du^1 du^2 \\ + (g_{12} b_{22} - g_{22} b_{12})(du^2)^2 = 0, \end{aligned}$$

to the finite element, the principal directions at the central node $(0, 0)$ are given by

$$\frac{du^2}{du^1} = \tan \theta = \frac{-B \pm \sqrt{B^2 - 4AC}}{2A}, \quad (27)$$

where

$$\begin{aligned} A &= 2bde - c(1 + e^2); \\ B &= 2b(1 + d^2) - 2a(1 + e^2); \\ C &= c(1 + d^2) - 2ade. \end{aligned} \quad (28)$$

Further details are contained in [53, Ch. 11].

ACKNOWLEDGMENT

M. Brady, S. Ullman, and T. Poggio provided support and participated in numerous discussions. Colleagues generously supplied test data for some of the experiments: P. Brou provided laser rangefinder data, E. Grimson, M. Kass, and K. Nishihara provided stereo data, K. Ikeuchi provided photometric stereo data, and J. Mahoney provided digital terrain maps.

REFERENCES

- [1] M. Abramowitz and I. A. Stegun, Eds., *Handbook of Mathematical Functions*. New York: Dover, 1965.
- [2] H. G. Barrow and J. M. Tenenbaum, "Recovering intrinsic scene characteristics from images," in *Computer Vision Systems*, A. Hanson and E. Riseman, Eds. New York: Academic, 1978, pp. 3-26.
- [3] —, "Reconstructing smooth surfaces from partial, noisy information," in *Proc. DARPA Image Understanding Workshop*, L. S. Baumann, Ed., Univ. Southern California, 1979, pp. 76-86.
- [4] —, "Interpreting line drawings as three-dimensional surfaces," *Artificial Intell.*, vol. 17, pp. 75-116, 1981.
- [5] P. J. Besl and R. C. Jain, "Invariant surface characteristics for three-dimensional object recognition in range images," *Comput. Vision, Graphics, Image Processing*, vol. 33, pp. 33-80, 1986.
- [6] A. Blake, "The least-disturbance principle and weak constraints," *tern Recognition Lett.*, vol. 1, pp. 393-399, 1983.
- [7] —, "Reconstructing a visible surface," in *Proc. Nat. Conf. AI (AAAI-84)*, Austin, TX, 1984, pp. 23-26.
- [8] A. Blake and A. Zisserman, "Invariant surface reconstruction using weak continuity constraints," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, Miami, FL, 1986, pp. 62-67.
- [9] G. Bolondi, F. Rocca, and S. Zanoletti, "Automatic contouring of faulted subsurfaces," *Geophys.*, vol. 41, pp. 1377-1393, 1976.
- [10] T. E. Boulton and J. R. Kender, "Visual surface reconstruction using sparse depth data," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, Miami, FL, 1986, pp. 68-76.
- [11] O. J. Braddick, F. W. Campbell, and J. Atkinson, "Channels in vision: Basic aspects," in *Handbook of Sensory Physiology: Perception*, vol. 8, R. Held, H. W. Leibowitz, and H. L. Teuber, Eds. Berlin: Springer, 1978, pp. 3-38.
- [12] J. M. Brady and B. K. P. Horn, "Rotationally symmetric operators for surface interpolation," *Comput. Vision, Graphics, Image Processing*, vol. 22, pp. 70-94, 1983.
- [13] J. M. Brady, J. Ponce, A. Yuille, and H. Asada, "Describing surfaces," *Comput. Vision, Graphics, Image Processing*, vol. 32, pp. 1-28, 1985.
- [14] I. C. Briggs, "Machine contouring using minimum curvature," *Geophys.*, vol. 39, pp. 39-48, 1974.
- [15] P. Burt and B. Julesz, "A disparity gradient limit for binocular fusion," *Science*, vol. 208, pp. 615-617, 1980.
- [16] G. Chavent, "Identification of distributed parameters," in *IFAC Symp. Identification and System Parameter Estimation*. New York: American Elsevier, 1973.
- [17] F. H. Clarke, *Optimization and Nonsmooth Analysis*. New York: Wiley-Interscience, 1983.
- [18] T. S. Collett, "Extrapolating and interpolating surfaces in depth," *Proc. Roy. Soc. London B*, vol. 224, pp. 43-56, 1985.
- [19] S. Cochran and M.-M. Medioni, "Implementation of a multiresolution surface reconstruction algorithm," *Dep. Elec. Eng. and Comput. Sci., Univ. Southern California, Los Angeles, Rep. ISG-108*, 1985.
- [20] M. P. do Carmo, *Differential Geometry of Curves and Surfaces*. Englewood Cliffs, NJ: Prentice-Hall, 1976.
- [21] Y. M. Ermoliev, "Methods of nondifferentiable and stochastic optimization and their applications," in *Progress in Nondifferentiable Optimization*, E. A. Nurminski, Ed., Int. Inst. Appl. Syst. Anal., Laxenburg, Austria, 1982.
- [22] T. J. Fan, G. Medioni, and R. Nevatia, "Description of surfaces from range data," in *Proc. DARPA Image Understanding Workshop*, Miami Beach, FL, L. S. Baumann, Ed., 1985, pp. 232-244.
- [23] R. Franke, "Thin plate splines with tension," *Comput. Aided Geometric Design*, vol. 2, pp. 87-95, 1985.
- [24] R. Franke and G. M. Neilson, "Surface approximation with imposed conditions," in *Surfaces in Computer Aided Geometric Design*, R. E. Barnhill and W. Boehm, Eds. Amsterdam, The Netherlands: North-Holland, 1983, pp. 135-146.
- [25] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-6, pp. 721-741, 1985.
- [26] J. J. Gibson, *The Perception of the Visual World*. Boston, MA: Houghton Mifflin, 1950.
- [27] W. E. L. Grimson, *From Images in Surfaces: A Computational Study of the Human Early Visual System*. Cambridge, MA: MIT Press, 1981.
- [28] —, "An implementation of a computational theory of visual surface interpolation," *Comput. Vision, Graphics, Image Processing*, vol. 22, pp. 39-69, 1983.
- [29] —, "Computational experiments with a feature based stereo algorithm," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-7, pp. 17-34, 1985.
- [30] W. E. L. Grimson and T. Pavlidis, "Discontinuity detection for visual surface reconstruction," *Comput. Vision, Graphics, Image Processing*, vol. 30, pp. 316-330, 1985.
- [31] W. Hackbusch, *Multi-Grid Methods and Applications*. Berlin: Springer-Verlag, 1985.
- [32] L. A. Hageman and D. M. Young, *Applied Iterative Methods*. New York: Academic, 1981.
- [33] J. G. Harris, "The coupled depth/slope approach to surface reconstruction," M.S. thesis, Dep. Elec. Eng. Comput. Sci., MIT, Cambridge, MA, 1986.
- [34] B. K. P. Horn, "Sequins and quills—A representation for surface topography," in *Representation of Three-Dimensional Objects*, R. Bajcsy, Ed. New York: Springer-Verlag, 1982.
- [35] W. J. Karplus, *Analog Simulation: Solution of Field Problems*. New York: McGraw-Hill, 1958.
- [36] M. Kass, "Computing visual correspondence," in *From Pixels to Predicates*, A. P. Pentland, Ed. Norwood, NJ: Ablex, 1986, pp. 78-92.
- [37] C. Koch, J. L. Marroquin, and A. L. Yuille, "Analog 'neuronal' networks in early vision," *Proc. Nat. Acad. Sci.*, p. 4263, 1986.
- [38] D. J. Langridge, "Detection of discontinuities in the first derivatives of surfaces," *Comput. Vision, Graphics, Image Processing*, vol. 27, pp. 291-308, 1984.
- [39] Y. Leclerc and S. W. Zucker, "The local structure of image discontinuities in one dimension," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-9, pp. 341-355, 1987.
- [40] D. Marr, *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco, CA: Freeman, 1982.
- [41] D. Marr and H. K. Nishihara, "Representation and recognition of the spatial organization of three-dimensional shapes," *Proc. Roy. Soc. London B*, vol. 200, 1978, pp. 269-294.
- [42] J. L. Marroquin, "Surface reconstruction preserving discontinuities," MIT A.I. Lab., Cambridge, MA, AI Memo 792, 1984.
- [43] D. Mumford and J. Shah, "Boundary detection by minimizing functionals, I," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, San Francisco, CA, 1985, pp. 22-26.
- [44] T. Poggio, V. Torre, and C. Koch, "Computational vision and regularization theory," *Nature*, vol. 317, pp. 314-319, 1985.
- [45] M. P. Polis and R. E. Goodson, "Parameter identification in distributed systems: A synthesizing overview," *Proc. IEEE*, vol. 64, pp. 45-61, 1976.
- [46] R. T. Rockafellar, *The Theory of Subgradients and its Application to Problems of Optimization: Convex and Nonconvex Functions*. Berlin: Helderman Verlag, 1981.
- [47] L. L. Schumaker, "Fitting surfaces to scattered data," in *Approximation Theory II*, G. G. Lorentz, C. K. Chui, and L. L. Schumaker, Eds. New York: Academic, 1976, pp. 203-267.
- [48] J. H. Shiau, "Smoothing spline estimation of functions with discontinuities," Ph.D. dissertation, Dep. Statistics, Univ. Wisconsin, Madison, 1985.
- [49] G. Strang and G. J. Fix, *An Analysis of the Finite Element Method*. Englewood Cliffs, NJ: Prentice-Hall, 1973.
- [50] R. Szilard, *Theory and Analysis of Plates: Classical and Numerical Methods*. Englewood Cliffs, NJ: Prentice-Hall, 1974.
- [51] D. Terzopoulos, "Multilevel computational processes for visual surface reconstruction," *Comput. Vision, Graphics, Image Processing*, vol. 24, pp. 52-96, 1983.
- [52] —, "The role of constraints and discontinuities in visible-surface reconstruction," in *Proc. 8th Int. Joint Conf. AI*, Karlsruhe, W. Germany, 1983, pp. 1073-1077.
- [53] —, "Multiresolution computation of visible surface representations," Ph.D. dissertation, Dep. Elec. Eng. Comput. Sci., MIT, Cambridge, MA, Jan. 1984.
- [54] —, "Concurrent multilevel relaxation," in *Proc. DARPA Image Understanding Workshop*, L. S. Baumann, Ed., Miami Beach, FL, 1985, pp. 156-162.
- [55] —, "Image analysis using multigrid relaxation methods," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-8, pp. 129-139, 1986.
- [56] —, "Regularization of inverse visual problems involving discontinuities," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-8, pp. 413-424, 1986.

- [57] A. N. Tikhonov and V. A. Arsenin, *Solutions of Ill-Posed Problems*. Washington, DC: Winston and Sons, 1977.
- [58] V. Torre and T. A. Poggio, "On edge detection," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-8, pp. 147-163, 1986.
- [59] B. C. Vemuri, A. Mitiche, and J. K. Aggarwal, "Curvature-based representation of objects from range data," *Image and Vision Computing*, vol. 4, pp. 107-114, 1986.
- [60] G. Wahba, "Spline bases, regularization, and generalized cross validation for solving approximation problems with large quantities of noisy data," in *Approximation Theory III*, E. W. Cheney, Ed. New York: Academic, 1980, pp. 905-912.
- [61] E. Wasserstrom, "Numerical solutions by the continuation method," *SIAM Rev.*, vol. 15, pp. 89-119, 1973.
- [62] R. J. Woodham, "Analyzing images of curved surfaces," *Artificial Intell.*, vol. 17, pp. 117-140, 1981.
- [63] O. C. Zienkiewicz, *The Finite Element Method*, 3rd ed. New York: McGraw-Hill, 1979.
- [64] S. W. Zucker, R. A. Hummel, and A. Rosenfeld, "An application of relaxation labeling to line and curve enhancement," *IEEE Trans. Comput.*, vol. C-26, pp. 394-403, 1977.
- [65] S. W. Zucker and P. Parent, "Multiple size operators and optimal curve finding," in *Multiresolution Image Processing and Analysis*, A. Rosenfeld, Ed. New York: Springer-Verlag, 1984, pp. 200-210.