The Topographic Primal Sketch

Robert M. Haralick
Layne T. Watson
Thomas J. Laffey
Departments of Computer Science and
Electrical Engineering
Virginia Polytechnic Institute and State University
Blacksburg, VA 24061

Abstract

A complete mathematical treatment is given for describing the topographic primal sketch of the underlying gray tone intensity surface of a digital image. Each picture element is independently classified and assigned a unique descriptive label, invariant under monotonically increasing gray tone transformations from the set (peak, pit, ridge, ravine, saddle, flat, and hillside), with hillside having subcategories (inflection point, slope, convex hill, concave hill, and saddle hill). The topographic classification is based on the first and second directional derivatives of the estimated image-intensity surface. A local, facet model, two-dimensional, cubic polynomial fit is done to estimate the image-intensity surface. Zero-crossings of the first directional derivative are identified as locations of interest in the image.

1. Introduction

Representing the fundamental structure of a digital image in a rich and robust way is a primary problem encountered in any general robotics computer-vision system that has to "understand" an image. The richness is needed so that shading, highlighting, and shadow information, which are usually present in real manufacturing assembly line situations, are encoded. Richness permits unambiguous object matching to be accomplished. Robustness is needed so that the representation is invariant with respect to monotonically increasing gray tone transformations.

This research has been supported by National Science Foundation grant MCS-8102872.

The International Journal of Robotics Research
Vol. 2, No. 1, Spring 1983
0278-3649/83/010050-23 $05.00/0

Current representations involving edges or the primal sketch as described by Marr (1976; 1980) are impoverished in the sense that they are insufficient for unambiguous matching. They also do not have the required invariance. Basic research is needed to (1) define an appropriate representation, (2) develop a theory that establishes its relationship to properties that three-dimensional objects manifest on the image, and (3) prove its utility in practice. Until this is done, computer-vision research must inevitably be more ad hoc sophistication than science.

The basis of the topographic primal sketch consists of the classification and grouping of the underlying image-intensity surface patches according to the categories defined by monotonic, gray tone, invariant functions of directional derivatives. Examples of such categories are peak, pit, ridge, ravine, saddle, flat, and hillside. From this initial classification, we can group categories to obtain a rich, hierarchical, and structurally complete representation of the fundamental image structure. We call this representation the *topographic primal sketch*.

Why do we believe that this topographic primal sketch can be the basis for computer vision? We believe it because the light-intensity variations on an image are caused by an object's surface orientation, its reflectance, and characteristics of its lighting source. If any of the three-dimensional intrinsic surface characteristics are to be detected, they will be detected owing to the nature of light-intensity variations. Thus, the first step is to discover a robust representation that can encode the nature of these light-intensity variations, a representation that does not change with strength of lighting or with gain settings on the sensing camera. The topographic classification does just that. The basic research issue is to define a set of categories sufficiently complete to form groupings and structures that have strong

relationships to the reflectances, surface orientations, and surface positions of the three-dimensional objects viewed in the image.

1.1. The Invariance Requirement

A digital image can be obtained with a variety of sensing-camera gain settings. It can be visually enhanced by an appropriate adjustment of the camera's dynamic range. The gain setting or the enhancing, point operator changes the image by some monotonically increasing function that is not necessarily linear. For example, nonlinear, enhancing, point operators of this type include histogram normalization and equal probability quantization.

In visual perception, exactly the same visual interpretation and understanding of a pictured scene occurs whether the camera's gain setting is low or high and whether the image is enhanced or unenhanced. The only difference is that the enhanced image has more contrast, is nicer to look at, and is understood more quickly by the human visual system.

This fact is important because it suggests that many of the current-low-level computer-vision techniques, which are based on edges, cannot ever hope to have the robustness associated with human visual perception. They cannot have the robustness, because they are inherently incapable of invariance under monotonic transformations. For example, edges based on zero-crossings of second derivatives will change in position as the monotonic gray tone transformation changes because convexity of a gray tone intensity surface is not preserved under such transformations. However, the topographic categories peak, pit, ridge, valley, saddle, flat, and hillside do have the required invariance.

1.2. Background

Marr (1976) argues that the first level of visual processing is the computation of a rich description of gray level changes present in an image, and that all subsequent computations are done in terms of this description, which he calls the *primal sketch*. Gray level changes are usually associated with edges, and Marr's primal sketch has, for each area of gray level change, a description that includes type, position, orientation, and fuzziness of edge. Marr (1980) illustrates that from this information it is sometimes possible to reconstruct the image to a reasonable degree. Unfortunately, as mentioned earlier, edge is not invariant with respect to monotonic image transformations; besides, it is not a rich enough structure. For example, difficulty has been experienced in using edges to accomplish unambiguous stereo matching.

The topographic primal sketch we are discussing as a basis for a representation has the required richness and invariance properties and is very much in the spirit of Marr's primal sketch and the thinking behind Ehrich's relational trees (Ehrich and Foith 1978). Instead of concentrating on gray level changes as edges as Marr does, or on one-dimensional extrema as Ehrich and Foith, we concentrate on all types of two-dimensional gray level variations. We consider each area on an image to be a spatial distribution of gray levels that constitutes a surface or facet of gray tone intensities having a specific surface shape. It is likely that, if we could describe the shape of the gray tone intensity surface for each pixel, then by assembling all the shape fragments we could reconstruct, in a relative way, the entire surface of the image's gray tone intensity values. The shapes that we already know about that have the invariance property are peak, pit, ridge, ravine, saddle, flat, and hillside, with hillside having noninvariant subcategories of slope, inflection, saddle hillside, convex hillside, and concave hillside.

Knowing that a pixel's surface has the shape of a peak does not tell us precisely where in the pixel the peak occurs; nor does it tell us the height of the peak or the magnitude of the slope around the peak. The topographic labeling, however, does satisfy Marr's (1976) primal sketch requirement in that it contains a symbolic description of the gray tone intensity changes. Furthermore, upon computing and binding to each topographic label numerical descriptors such as gradient magnitude and direction, directions of the extrema of the second directional derivative along with their values, a reasonable absolute description of each surface shape can be obtained.

1.3. Facet Model

The *facet model* states that all processing of digital image data has its final authoritative interpretation relative to what the processing does to the underlying gray tone intensity surface. The digital image's pixel values are noisy sampled observations of the underlying surface. Thus, in order to do any processing, we at least have to estimate at each pixel position what this underlying surface is. This requires a model that describes what the general form of the surface would be in the neighborhood of any pixel if there were no noise. To estimate the surface from the neighborhood around a pixel then amounts to estimating the free parameters of the general form. It is important to note that if a different general form is assumed, then a different estimate of the surface is produced. Thus the assumption of a particular general form is necessary and has consequences.

The general form we use is a bivariate cubic. We assume that the neighborhood around each pixel is suitably fit by a bivariate cubic (Haralick 1981; 1982). Having estimated this surface around each pixel, the first and second directional derivatives are easily computed by analytic means. The topographic classification of the surface facet is based totally on the first and second directional derivatives. We classify each surface point as peak, pit, ridge, ravine, saddle, flat, or hillside, with hillside being broken down further into the subcategories inflection point, convex hill, concave hill, saddle hill, and slope. Our set of topographic labels is complete in the sense that every combination of values of the first and second directional derivative is uniquely assigned to one of the classes.

1.4. Previous Work

Detection of topographic structures in a digital image is not a new idea. There has been a wide variety of techniques described to detect pits, peaks, ridges, ravines, and the like.

Peuker and Johnston (1972) characterize the surface shape by the sequence of positive and negative differences as successive surrounding points are compared to the central point. Peuker and Douglas (1975) describe several variations of this method for detecting one of the shapes from the set (pit, peak, pass, ridge, ravine, break, slope, flat). They start with the most frequent feature (slope) and proceed to the less frequent, thus making it an order-dependent algorithm.

Johnston and Rosenfeld (1975) attempt to find peaks by finding all points P such that no points in an n-by-n neighborhood surrounding P have greater elevation than P. Pits are found in an analogous manner. To find ridges, they identify points that are either east-west or north-south elevation maxima. This is done using a "smoothed" array in which each point is given the highest elevation in a 2×2 square containing it. East-west and north-south maxima are also found on this array. Ravines are found in a similar manner.

Paton (1975) uses a six-term quadratic expansion in Legendre polynomials fitted to a small disk around each pixel. The most significant coefficients of the second-order polynomial yield a descriptive label chosen from the set (constant, ridge, valley, peak, bowl, saddle, ambiguous). He uses the continuous least-squares-fit formulation in setting up the surface-fit equations as opposed to the discrete least-squares fit used in the facet model. The continuous fit is a more expensive computation than the discrete fit and results in a steplike approximation.

Grender's (1976) algorithm compares the gray level elevation of a central point with surrounding elevations at a given distance around the perimeter of a circular window; the radius of the window may be increased in successive passes through the image. His topographic labeling set consists of slope, ridge, valley, knob, sink, saddle.

Toriwaki and Fukumara (1978) take a totally different approach from all the others. They use two local features of gray level pictures, connectivity number, and coefficient of curvature for classification of the pixel into peak, pit, ridge, ravine, hillside, pass. They then describe how to extract structural information from the image once the labelings have been made. This structural information consists of ridge-lines, ravine-lines, and the like.

Hsu, Mundy, and Beaudet (1978) use a quadratic surface approximation at every point on the image

surface. The principal axes of the quadratic approximation are used as directions in which to segment the image. Lines emanating from the center pixel in these directions thus provide natural boundaries of patches approximating the surface. The authors then selectively generate the principal axes from some critical points distributed over an image and interconnect them into a network to get an approximation of the image data. In this network, which they call the *web representation,* the axes divide the image into regions and show important features such as edges and peaks. They are then able to extract a set of primitive features from the nodes of the network by mask matching. Global features, such as ridge-lines, are obtained by state transition rules.

Lee and Fu (1981) define a set of 3 × 3 templates that they convolve over the image to give each class except plain a *figure of merit.* Their set of labels includes none, plain, slope, ridge, valley, foot, shoulder. Thresholds are used to determine into which class the pixel will fall. In their scheme, a pixel may satisfy the definition of zero, one, or more than one class. Ambiguity is resolved by choosing the class with the highest figure of merit.

1.5. A Mathematical Approach

From the previous discussion, one can see that a wide variety of methods and labels has been proposed to describe the topographic structure in a digital image. Some of the methods require multiple passes through the image, while others may give ambiguous labels to a pixel. Many of the methods are heuristic in nature. The Hsu, Mundy, and Beudet (1978) approach is the most similar to the one discussed here.

Our classification approach is based on the estimation of the first- and second-order directional derivatives. Thus, we regard the digital-picture function as a sampling of the underlying function f, where some kind of random noise is added to the true function values. To estimate the first and second partials, we must assume some kind of parametric form for the underlying function f. The classifier must use the sampled brightness values of the digital-picture function to estimate the parameters and then make decisions regarding the locations of relative extrema of partial derivatives based on the estimated values of the parameters.

In Section 2, we will discuss the mathematical properties of the topographic structures in terms of the directional derivatives in the continuous surface domain. Because a digital image is a sampled surface and each pixel has an area associated with it, characteristic topographic structures may occur anywhere within a pixel's area. Thus, the implementation of the mathematical topographic definitions is not entirely trivial.

In Section 3 we will discuss the implementation of the classification scheme on a digital image. To identify categories that are local one-dimensional extrema, such as peak, pit, ridge, ravine, and saddle, we search inside the pixel's area for a zero-crossing of the first directional derivative. The directions in which we seek the zero-crossing are along the lines of extreme curvature.

In Section 4, we will discuss the local cubic estimation scheme. In Section 5, we will summarize the algorithm for topographic classification using the local facet model. In Section 6, we will show the results of the classifier on several test images.

2. The Mathematical Classification of Topographic Structures

In this section, we formulate our notion of topographic structures on continuous surfaces and show their invariance under monotonically increasing gray tone transformations. In order to understand the mathematical properties used to define our topographic structures, one must understand the idea of the *directional derivative* discussed in most advanced calculus books. For completeness, we first give the definition of the directional derivative, then the definitions of the topographic labels. Finally, we show the invariance under monotonically increasing gray tone transformations.

2.1. The Directional Derivative

In two dimensions, the rate of change of a function f depends on direction. We denote the directional

derivative of f at the point (r, c) in the direction β by $f'_\beta(r, c)$. It is defined as

$$f'_\beta(r, c) = \lim_{h \to 0} \frac{f(r + h^* \sin \beta, c + h^* \cos \beta) - f(r, c)}{h}.$$

The direction angle β is the clockwise angle from the column axis. It follows directly from this definition that

$$f'_\beta(r, c) = \frac{\partial f}{\partial r}(r, c) * \sin \beta + \frac{\partial f}{\partial c}(r, c) * \cos \beta.$$

We denote the second derivative of f at the point (r, c) in the direction β by $f''_\beta(r, c)$, and it follows that

$$f''_\beta = \frac{\partial^2 f}{\partial r^2} * \sin^2 \beta + 2 * \frac{\partial^2 f}{\partial r\, \partial c} * \sin \beta * \cos \beta + \frac{\partial^2 f}{\partial c^2} * \cos^2 \beta.$$

The *gradient* of f is a vector whose magnitude,

$$\left(\left(\frac{\partial f}{\partial r}\right)^2 + \left(\frac{\partial f}{\partial c}\right)^2\right)^{1/2}$$

at a given point (r, c) is the maximum rate of change of f at that point, and whose direction,

$$\tan^{-1}\left(\frac{\frac{\partial f}{\partial r}}{\frac{\partial f}{\partial c}}\right)$$

is the direction in which the surface has the greatest rate of change.

2.2. The Mathematical Properties

We will use the following notation to describe the mathematical properties of our various topographic categories for continuous surfaces. Let

- ∇f = gradient vector of a function f;
- $\|\nabla f\|$ = gradient magnitude;
- $\omega^{(1)}$ = unit vector in direction in which second directional derivative has greatest magnitude;
- $\omega^{(2)}$ = unit vector orthogonal to $\omega^{(1)}$;
- λ_1 = value of second directional derivative in the direction of $\omega^{(1)}$;
- λ_2 = value of second directional derivative in the direction of $\omega^{(2)}$;
- $\nabla f \cdot \omega^{(1)}$ = value of first directional derivative in the direction of $\omega^{(1)}$; and
- $\nabla f \cdot \omega^{(2)}$ = value of first directional derivative in the direction of $\omega^{(2)}$.

Without loss of generality, we assume $|\lambda_1| >= |\lambda_2|$.

Each type of topographic structure in our classification scheme is defined in terms of the above quantities. In order to calculate these values, the first- and second-order partials with respect to r and c need to be approximated. These five partials are as follows:

$$\partial f/\partial r,\ \partial f/\partial c,\ \partial^2 f/\partial r^2,\ \partial^2 f/\partial c^2,\ \partial^2 f/\partial r\, \partial c.$$

The gradient vector is simply $(\partial f/\partial r,\ \partial f/\partial c)$. The second directional derivatives may be calculated by forming the *Hessian* where the Hessian is a 2×2 matrix defined as

$$H = \begin{vmatrix} \partial^2 f/\partial r^2 & \partial^2 f/\partial r\, \partial c \\ \partial^2 f/\partial c\, \partial r & \partial^2 f/\partial c^2 \end{vmatrix}.$$

Hessian matrices are used extensively in nonlinear programming. Only three parameters are required to determine the Hessian matrix H, since the order of differentiation of the cross partials may be interchanged. That is,

$$\partial^2 f/\partial r\, \partial c = \partial^2 f/\partial c\, \partial r.$$

The eigenvalues of the Hessian are the values of the extrema of the second directional derivative, and their associated eigenvectors are the directions in which the second directional derivative is extremized. This can easily be seen by rewriting f''_β as the quadratic form

$$f''_\beta = (\sin \beta \cos \beta) * H * \begin{vmatrix} \sin \beta \\ \cos \beta \end{vmatrix}.$$

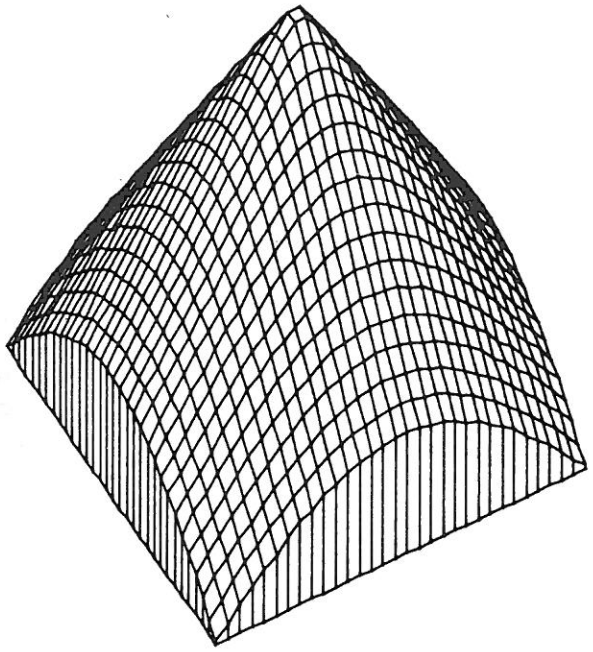

Fig. 1. Right circular cone.

Thus

$$H\omega^{(1)} = \lambda_1\omega^{(1)}, \text{ and } H\omega^{(2)} = \lambda_2\omega^{(2)}.$$

Furthermore, the two directions represented by the eigenvectors are orthogonal to one another. Since H is a 2×2 symmetric matrix, calculation of the eigenvalues and eigenvectors can be done efficiently and accurately using the method of Rutishauser (1971). We may obtain the values of the first directional derivative by simply taking the dot product of the gradient with the appropriate eigenvector:

$$\nabla f \cdot \omega^{(1)}$$
$$\nabla f \cdot \omega^{(2)}.$$

There is a direct relationship between the eigenvalues λ_1 and λ_2 and curvature in the directions $\omega^{(1)}$ and $\omega^{(2)}$: When the first directional derivative $\nabla f \cdot \omega^{(1)} = 0$, then $\lambda_1/(1 + (\nabla f \cdot \nabla f))^{1/2}$ is the curvature in the direction $\omega^{(1)}$, $i = 1$ or 2.

Having the gradient magnitude and direction and the eigenvalues and eigenvectors of the Hessian, we can describe the topographic classification scheme.

2.2.1. *Peak*

A peak (knob) occurs where there is a local maxima in all directions. In other words, we are on a peak if, no matter what direction we look in, we see no point that is as high as the one we are on (Fig. 1). The curvature is downward in all directions. At a peak the gradient is zero, and the second directional derectional derivative is negative in all directions. To test whether the second directional derivative is negative in all directions, we just have to examine the value of the second directional derivative in the directions that make it smallest and largest. A point is therefore classified as a peak if it satisfies the following conditions:

$$\|\nabla f\| = 0, \lambda_1 < 0, \lambda_2 < 0.$$

2.2.2. *Pit*

A pit (sink, bowl) is identical to a peak except that it is a local minima in all directions rather than a local maxima. At a pit the gradient is zero, and the second directional derivative is positive in all directions. A point is classified as a pit if it satisfies the following conditions:

$$\|\nabla f\| = 0, \lambda_1 > 0, \lambda_2 > 0.$$

2.2.3. *Ridge*

A ridge occurs on a ridge-line, a curve consisting of a series of ridge points. As we walk along the ridge-line, the points to the right and left of us are lower than the ones we are on. Furthermore, the ridge-line may be flat, slope upward, slope downward, curve upward, or curve downward. A ridge occurs where there is a local maximum in one direction, as illustrated in Fig. 2. Therefore, it must have negative second-directional derivative in the direction across the ridge and also a zero first-directional derivative in that same direction. The direction in which the local maximum occurs may correspond to either of the directions in which the curvature is "extremized," since the ridge itself may be curved. For nonflat ridges, this leads to the first two cases below for ridge characterization. If the ridge is flat,

Fig. 2. Saddle surface.

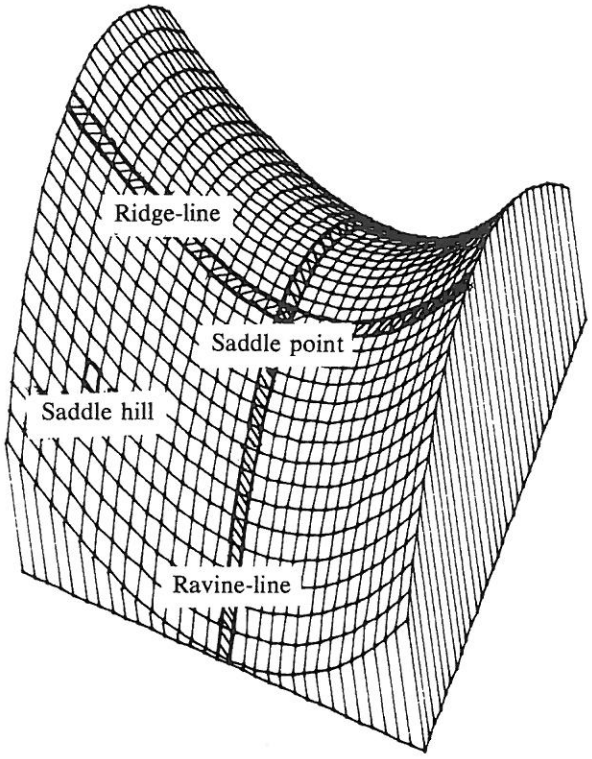

then the ridge-line is horizontal and the gradient is zero along it. This corresponds to the third case. The defining characteristic is that the second directional derivative in the direction of the ridge-line is zero, while the second directional derivative across the ridge-line is negative. A point is therefore classified as a ridge if it satisfies any one of the following three sets of conditions:

$$\|\nabla f\| \neq 0,\ \lambda_1 < 0,\ \nabla f \cdot \omega^{(1)} = 0$$

or

$$\|\nabla f\| \neq 0,\ \lambda_2 < 0,\ \nabla f \cdot \omega^{(2)} = 0$$

or

$$\|\nabla f\| = 0,\ \lambda_1 < 0,\ \lambda_2 = 0.$$

A geometric way of thinking about the definition for ridge is to realize that the condition $\nabla f \cdot \omega^{(1)} = 0$ means that the gradient direction (which is defined for nonzero gradients) is orthogonal to the direction $\omega^{(1)}$ of extremized curvature.

2.2.4. *Ravine*

A ravine (valley) is identical to a ridge except that it is a local minimum (rather than maximum) in one direction. As we walk along the ravine-line, the points to the right and left of us are higher than the one we are on (see Fig. 2). A point is classified as a ravine if it satisfies any one of the following three sets of conditions:

$$\|\nabla f\| \neq 0,\ \lambda_1 > 0,\ \nabla f \cdot \omega^{(1)} = 0$$

or

$$\|\nabla f\| \neq 0,\ \lambda_2 > 0,\ \nabla f \cdot \omega^{(2)} = 0$$

or

$$\|\nabla f\| = 0,\ \lambda_1 > 0,\ \lambda_2 = 0.$$

2.2.5. *Saddle*

A saddle occurs where there is a local maximum in one direction and a local minimum in a perpendicular direction, as illustrated in Fig. 2. A saddle must therefore have positive curvature in one direction and negative curvature in a perpendicular direction. At a saddle, the gradient magnitude must be zero and the extrema of the second directional derivative must have opposite signs. A point is classified as a saddle if it satisfies the following conditions:

$$\|\nabla f\| = 0,\ \lambda_1 * \lambda_2 < 0.$$

2.2.6. *Flat*

A flat (plain) is a simple, horizontal surface, as illustrated in Fig. 3. It, therefore, must have zero gradient and no curvature. A point is classified as a flat if it satisfies the following conditions:

$$\|\nabla f\| = 0,\ \lambda_1 = 0,\ \lambda_2 = 0.$$

Fig. 3. Hillside.

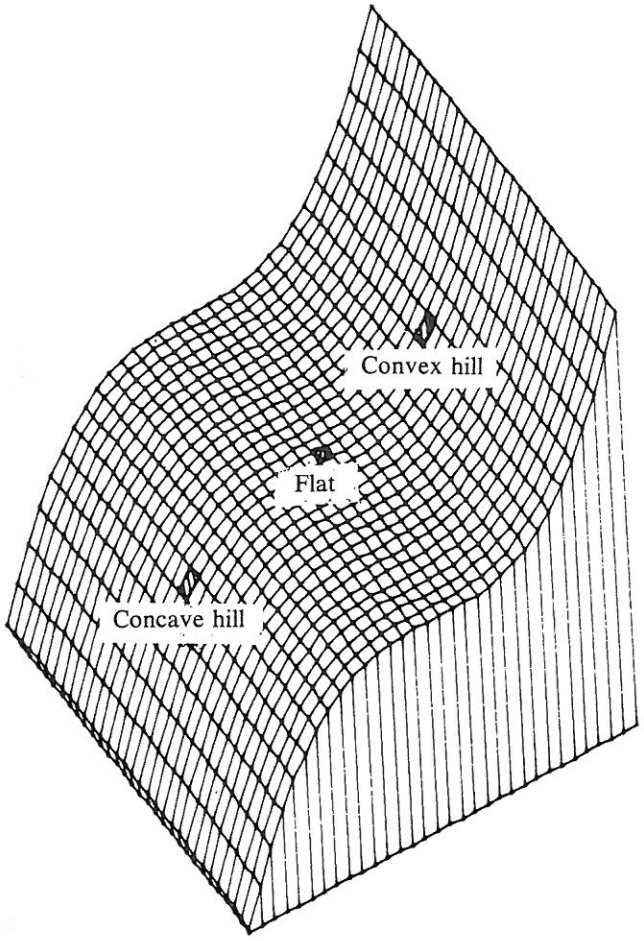

Given that the above conditions are true, a flat may be further classified as a *foot* or *shoulder*. A foot occurs at that point where the flat just begins to turn up into a hill. At this point, the third directional derivative in the direction toward the hill will be nonzero, and the surface increases in this direction. The shoulder is an analogous case and occurs where the flat is ending and turning down into a hill. At this point, the maximum magnitude of the third directional derivative is nonzero, and the surface decreases in the direction toward the hill. If the third directional derivative is zero in all directions, then we are in a flat, not near a hill. Thus a flat may be further qualified as being a foot or shoulder, or not qualified at all.

2.2.7. *Hillside*

A hillside point is anything not covered by the previous categories. It has a nonzero gradient and no strict extrema in the directions of maximum and minimum second directional derivative. If the hill is simply a tilted flat (i.e., has constant gradient), we call it a *slope*. If its curvature is positive (upward), we call it a *convex hill*. If its curvature is negative (downward), we call it a *concave hill*. If the curvature is up in one direction and down in a perpendicular direction, we call it a *saddle hill*. A saddle hill is illustrated in Fig. 2, and the slope, convex hill, and concave hill are illustrated in Fig. 3.

A point on a hillside is an *inflection point* if it has a zero-crossing of the second directional derivative taken in the direction of the gradient. The inflection-point class is the same as the *step edge* defined by Haralick (1982), who classifies a pixel as a step edge if there is some point in the pixel's area having a zero-crossing of the second directional derivative taken in the direction of the gradient.

To determine whether a point is a hillside, we just take the complement of the disjunction of the conditions given for all the previous classes. Thus if there is no curvature, then the gradient must be nonzero. If there is curvature, then the point must not be a relative extremum. Therefore, a point is classified as a hillside if all three sets of the following conditions are true ('$\rightarrow$' represents the operation of logical implication):

$$\lambda_1 = \lambda_2 = 0 \rightarrow \|\nabla f\| \neq 0,$$

and

$$\lambda_1 \neq 0 \rightarrow \nabla f \cdot \omega^{(1)} \neq 0,$$

and

$$\lambda_2 \neq 0 \rightarrow \nabla f \cdot \omega^{(2)} \neq 0.$$

Rewritten as a disjunction of clauses rather than a conjunction of clauses, a point is classified as a hillside if any one of the following four sets of conditions are true:

$$\nabla f \cdot \omega^{(1)} \neq 0, \nabla f \cdot \omega^{(2)} \neq 0$$

or

$$\nabla f \cdot \omega^{(1)} \neq 0,\ \lambda_2 = 0$$

or

$$\nabla f \cdot \omega^{(2)} \neq 0,\ \lambda_1 = 0$$

or

$$\|\nabla f\| \neq 0,\ \lambda_1 = 0,\ \lambda_2 = 0.$$

We can differentiate between different classes of hillsides by the values of the second directional derivative. The distinction can be made as follows:

Slope	if $\lambda_1 = \lambda_2 = 0$
Convex	if $\lambda_1 > = \lambda_2 > = 0,\ \lambda_1 \neq 0$
Concave	if $\lambda_1 < = \lambda_2 < = 0,\ \lambda_1 \neq 0$
Saddle hill	if $\lambda_1 * \lambda_2 < 0$

A slope, convex, concave, or saddle hill is classified as an inflection point if there is a zero-crossing of the second directional derivative in the direction of maximum first directional derivative (i.e., the gradient).

2.2.8. *Summary of the Topographic Categories*

A summary of the mathematical properties of our topographic structures on continuous surfaces can be found in Table 1. The table exhaustively defines the topographic classes by their gradient magnitude, second directional derivative extrema values, and the first directional derivatives taken in the directions which extremize second directional derivatives. Each entry in the table is either 0, +, −, or ∗. The 0 means not significantly different from zero; + means significantly different from zero on the positive side; − means significantly different from zero on the negative side, and ∗ means it does not matter. The label "Cannot occur" means that it is impossible for the gradient to be nonzero and the first directional derivative to be zero in two orthogonal directions.

From the table, one can see that our classification scheme is complete. All possible combinations of first and second directional derivatives have a corresponding entry in the table. Each topographic category has a set of mathematical properties that uniquely determines it.

Table 1. Mathematical Properties of Topographic Structures

$\|\nabla f\|$	λ_1	λ_2	$\nabla f \cdot \omega^{(1)}$	$\nabla f \cdot \omega^{(2)}$	Label
0	−	−	0	0	Peak
0	−	0	0	0	Ridge
0	−	+	0	0	Saddle
0	0	0	0	0	Flat
0	+	−	0	0	Saddle
0	+	0	0	0	Ravine
0	+	+	0	0	Pit
+	−	−	−, +	−, +	Hillside
+	−	∗	0	∗	Ridge
+	∗	−	∗	0	Ridge
+	−	0	−, +	∗	Hillside
+	−	+	−, +	−, +	Hillside
+	0	0	∗	∗	Hillside
+	+	−	−, +	−, +	Hillside
+	+	0	−, +	∗	Hillside
+	+	∗	0	∗	Ravine
+	∗	+	∗	0	Ravine
+	+	+	−, +	−, +	Hillside
+	∗	∗	0	0	Cannot occur

(Note: Special attention is required for the degenerate case $\lambda_1 = \lambda_2 \neq 0$, which implies that $\omega^{(1)}$ and $\omega^{(2)}$ can be *any* two orthogonal directions. In this case, there *always* exists an extreme direction ω which is orthogonal to ∇f, and thus the first directional derivative $\nabla f \cdot \omega$ is *always* zero in an extreme direction. To avoid spurious zero directional derivatives, we choose $\omega^{(1)}$ and $\omega^{(2)}$ such that $\nabla f \cdot \omega^{(1)} \neq 0$ and $\nabla f \cdot \omega^{(2)} \neq 0$, unless the gradient is zero.)

2.3. The Invariance of the Topographic Categories

In this section, we show that the topographic labels (peak, pit, ridge, ravine, saddle, flat, and hillside), the gradient direction, and directions of second directional derivative extrema for peak, pit, ridge, ravine, and saddle are all invariant under monotoni-

cally increasing gray tone transformations. We take *monotonically increasing* to mean positive derivative everywhere.

Let the original underlying gray tone surface be $f(r, c)$. Let w be a monotonically increasing gray tone transformation, and let $g(r, c)$ denote the transformed image: $g(r, c) = w(f(r, c))$. It is directly derivable that

$$g'_\beta(r, c) = w'(f(r, c)) * f'_\beta(r, c),$$

from which we obtain that

$$g''_\beta(r, c) = w'(f(r, c)) * f''_\beta(r, c) + w''(f_\beta(r, c)) * f'_\beta(r, c)^2.$$

Let us fix a position (r, c). Since w is a monotonically increasing function, w' is positive. In particular, w' is not zero. Hence the direction β which maximizes g'_β also maximizes f'_β, thereby showing that the gradient directions are the same. The categories peak, pit, ridge, ravine, saddle, and flat all have in common the essential property that the first directional derivative is zero when taken in a direction that extremizes the second directional derivative. To see the invariance, let β be an extremizing direction of f''_β. Then for points (r, c) having a label (peak, pit, ridge, ravine, saddle, or flat), $f'_\beta(r, c) = 0$, and $\partial f''_\beta(r, c)/\partial\beta = 0$. Notice that

$$\frac{\partial g''_\beta}{\partial \beta} = w' * \frac{\partial}{\partial \beta} f''_\beta + 2 * w' * f'_\beta \frac{\partial f'_\beta}{\partial \beta} + (f'_\beta)^2 \frac{\partial w''}{\partial \beta}.$$

Hence for these points, $g'_\beta(r, c) = 0$, and

$$\partial g''_\beta(r, c)/\partial\beta = 0,$$

thereby showing that at these points the directions that extremize f''_β are precisely the directions that extremize g''_β, and that g''_β will always have the same sign as f''_β. A similar argument shows that if β extremizes g''_β and satisfies $g'_\beta = 0$, then β must also extremize f''_β and satisfy $f'_\beta = 0$. Therefore, any points in the original image with the labels peak, pit, ridge, saddle, or flat retain the same label in the transformed image and, conversely, any points in the transformed image will have the same label in the original image.

Any pixel with a label not in the set (peak, pit, ridge, ravine, saddle, and flat) must have a hillside label. Thus, a point labeled hillside must be transformed to a hillside-labeled point. However, the subcategories (inflection point, slope, convex hill, concave hill, and saddle hill) may change under the gray tone transformation.

2.4. Ridge and Ravine Continua

Although the definitions given for ridge and ravine are intuitively pleasing, they may lead to the unexpected consequence of having entire areas of a surface classified as all ridge or all ravine. To see how this can occur, observe that the eigenvalue $\lambda = \lambda(r, c)$ satisfies

$$\lambda(r, c) = \frac{1}{2}\left|\frac{\partial^2 f}{\partial r^2}(r, c) + \frac{\partial^2 f}{\partial c^2}(r, c)\right| \pm \left|\left|\frac{\partial^2 f}{\partial r\, \partial c}(r, c)\right|^2 + \left|\frac{1}{2}\left|\frac{\partial^2 f}{\partial r^2}(r, c) - \frac{\partial^2 f}{\partial c^2}(r, c)\right|\right|^2\right|^{1/2}.$$

For there to be a ridge or ravine at a point (r, c), the corresponding eigenvector $\omega(r, c)$ must be perpendicular to the gradient direction. Therefore, $\nabla f \cdot \omega = 0$. If this equation holds for a point (r, c) and not all points in a small neighborhood about (r, c), there is a ridge or ravine in the commonly understood sense. However, if this equation holds for all points in a neighborhood about (r, c), then we have a ridge or ravine continuum by the criteria of Sections 2.2.3 and 2.2.4.

Unfortunately, there are "nonpathologic" surfaces having ridge or ravine continuums. Simple, radially symmetric examples include the inverted right circular cone defined by

$$f(r, c) = (r^2 + c^2)^{1/2},$$

the hemisphere defined by

$$f(r, c) = (k^2 - r^2 - c^2)^{1/2},$$

or, in fact, any function of the form $h(r^2 + c^2)$. In

the case of the cone, the gradient is proportional to (r, c), and the unnormalized eigenvectors corresponding to eigenvalues

$$\lambda(r, c) = (r^2 + c^2)^{-1/2} \text{ and } 0$$

are $(-c, r)$ and (r, c) respectively. The eigenvector corresponding to the nonzero eigenvalue is orthogonal to the gradient direction. The entire surface of the inverted cone, except for the apex, is a ravine. Other, nonradially symmetric examples exist as well.

The identification of points that are really ridge or ravine continuums can be made as a postprocessing step. Points that are labeled as ridge or ravine and that have neighboring points in a direction orthogonal to the gradient that are also labeled ridge or ravine are ridge or ravine continuums. These continuums can be reclassified as hillsides.

3. The Topographic Classification Algorithm

The definitions of Section 2 cannot be used directly since there is a problem of where in a pixel's area to apply the classification. If the classification were only applied to the point at the center of each pixel, then a pixel having a peak near one of its corners, for example, would get classified as a concave hill rather than as a peak. The problem is that the topographic classification we are interested in must be a sampling of the actual topographic surface classes. Most likely, the interesting categories of peak, pit, ridge, ravine, and saddle will never occur precisely at a pixel's center, and if they do occur in a pixel's area, then the pixel must carry that label rather than the class label of the pixel's center point. Thus one problem we must solve is to determine the dominant label for a pixel given the topographic class label of every point in the pixel. The next problem we must solve is to determine, in effect, the set of all topographic classes occurring within a pixel's area without having to do the impossible brute-force computation.

For the purpose of solving these problems, we divide the set of topographic labels into two subsets: (1) those that indicate that a strict, local, one-dimensional extremum has occurred (peak, pit, ridge, ravine, and saddle) and (2) those that do not indicate that a strict, local, one-dimensional extremum has occurred (flat and hillside). By *one-dimensional,* we mean along a line (in a particular direction). A strict, local, one-dimensional extremum can be located by finding those points within a pixel's area where a zero-crossing of the first directional derivative occurs.

So that we do not search the pixel's entire area for the zero-crossing, we only search in the directions of extreme second directional derivative, $\omega^{(1)}$ and $\omega^{(2)}$. Since these directions are well aligned with curvature properties, the chance of overlooking an important topographic structure is minimized, and, more importantly, the computational cost is small.

When $\lambda_1 = \lambda_2 \neq 0$, the directions $\omega^{(1)}$ and $\omega^{(2)}$ are not uniquely defined. We handle this case by searching for a zero-crossing in the direction given by $\mathbf{H}^{-1} * \nabla f$. This is the *Newton direction,* and it points directly toward the extremum of a quadratic surface.

For inflection-point location (first derivative extremum), we search along the gradient direction for a zero-crossing of second directional derivative. For one-dimensional extrema, there are four cases to consider: (1) no zero-crossing, (2) one zero-crossing, (3) two zero-crossings, and (4) more than two zero-crossings. The next four sections discuss these cases.

3.1. Case One: No Zero-Crossing

If no zero-crossing is found along either of the two extreme directions within the pixel's area, then the pixel cannot be a local extremum and therefore must be assigned a label from the set (flat or hillside). If the gradient is zero, we have a flat. If it is nonzero, we have a hillside. If the pixel is a hillside, we classify it further into inflection point, slope, convex hill, concave hill, or saddle hill. If there is a zero-crossing of the second directional derivative in the direction of the gradient within the pixel's area, the pixel is classified as an inflection point. If no such zero-crossing occurs, the label assigned to the pixel is based on the gradient magnitude and Hessian eigenvalues calculated at the center of the pixel, local coordinates (0, 0), as in Table 2.

Table 2. Pixel Label Calculation for Case One: No Zero-Crossing

$\|\nabla f\|$	λ_1	λ_2	Label
0	0	0	Flat
+	−	−	Concave hill
+	−	0	Concave hill
+	−	+	Saddle hill
+	0	0	Slope
+	+	−	Saddle hill
+	+	0	Convex hill
+	+	+	Convex hill

3.2. Case Two: One Zero-Crossing

If a zero-crossing of the first directional derivative is found within the pixel's area, then the pixel is a strict, local, one-dimensional extremum and must be assigned a label from the set (peak, pit, ridge, ravine, or saddle). At the location of the zero-crossing, the Hessian and gradient are recomputed, and if the gradient magnitude at the zero-crossing is zero, Table 3 is used.

If the gradient magnitude is nonzero, then the choice is either ridge or ravine. If the second directional derivative in the direction of the zero-crossing is negative, we have a ridge. If it is positive, we have a ravine. If it is zero, we compare the function value at the center of the pixel, $f(0, 0)$, with the function value at the zero-crossing, $f(r, c)$. If $f(r, c)$ is greater than $f(0, 0)$, we call it a ridge, otherwise we call it a ravine.

3.3. Case Three: Two Zero-Crossings

If we have two zero-crossings of the first directional derivative, one in each direction of extreme curvature, then the Hessian and gradient must be recomputed at each zero-crossing. Using the procedure described in Section 3.2, we assign a label to each zero-crossing. We call these labels LABEL1 and LABEL2. The final classification given the pixel is based on these two labels and is given in Table 4.

If both labels are identical, the pixel is given that label. In the case of both labels being ridge, the pixel

Table 3. Pixel Label Calculation for Case Two: One Zero-Crossing

$\|\nabla f\|$	λ_1	λ_2	Label
0	−	−	Peak
0	−	0	Ridge
0	−	+	Saddle
0	+	−	Saddle
0	+	0	Ravine
0	+	+	Pit

Table 4. Final Pixel Classification, Case Three: Two Zero-Crossings

LABEL1	LABEL2	Resulting Label
Peak	Peak	Peak
Peak	Ridge	Peak
Pit	Pit	Pit
Pit	Ravine	Pit
Saddle	Saddle	Saddle
Ridge	Ridge	Ridge
Ridge	Ravine	Saddle
Ridge	Saddle	Saddle
Ravine	Ravine	Ravine
Ravine	Saddle	Saddle

may actually be a peak, but experiments have shown that this case is rare. An anlogous argument can be made for both labels being ravine. If one label is ridge and the other ravine, this indicates we are at or very close to a saddle point, and thus the pixel is classified as a saddle. If one label is peak and the other ridge, we choose the category giving us the "most information," which in this case is peak. The peak is a local maximum in all directions, while the ridge is a local maximum in only one direction. Thus, peak conveys more information about the image surface. An analogous argument can be made if the labels are pit and ravine. Similarly, a saddle gives us more information than a ridge or valley. Thus, a pixel is assigned saddle if its zero-crossings have been labeled ridge and saddle or ravine and saddle.

It is apparent from Table 4 that not all possible label combinations are accounted for. Some combi-

nations, such as peak and pit, are omitted because of the assumption that the underlying surface is smooth and sampled frequently enough that a peak and pit will not both occur within the same pixel's area. If such a case occurs, our convention is to choose arbitrarily one of LABEL1 or LABEL2 as the resulting label for the pixel.

3.4. Case Four: More than Two Zero-Crossings

If more than two zero-crossings occur within a pixel's area, then in at least one of the extrema directions there are two zero-crossings. If this happens, we choose the zero-crossing closest to the pixel's center and ignore the other. If we ignore the further zero-crossings, then this case is identical to case 3. This situation has yet to occur in our experiments.

4. Surface Estimation

In this section we discuss the estimation of the parameters required by the topographic classification scheme of Section 2 using the local cubic facet model (Haralick 1981). It is important to note that the classification scheme of Section 2 and the algorithm of Section 3 are independent of the method used to estimate the first- and second-order partials of the underlying digital image-intensity surface at each sampled point. Although we are currently using the cubic model and discuss it here, we expect that a spline-based estimation scheme or a discrete-cosines estimation scheme may, in fact, provide better estimates.

4.1. Local Cubic Facet Model

In order to estimate the required partial derivatives, we perform a least-squares fit with a two-dimensional surface, f, to a neighborhood of each pixel. It is required that the function f be continuous and have continuous first- and second-order partial derivatives with respect to r and c in a neighborhood around each pixel in the rc plane.

We choose f to be a cubic polynomial in r and c expressed as a combination of discrete orthogonal polynomials. The function f is the best discrete least-squares polynomial approximation to the image data in each pixel's neighborhood. More details can be found in Haralick's paper (1981), in which each coefficient of the cubic polynomial is evaluated as a linear combination of the pixels in the fitting neighborhood.

To express the procedure precisely and without reference to a particular set of polynomials tied to neighborhood size, we will canonically write the fitted bicubic surface for each fitting neighborhood as

$$\begin{aligned} f(r, c) = {} & k_1 + k_2 r + k_3 c \\ & + k_4 r^2 + k_5 rc + k_6 c^2 \\ & + k_7 r^3 + k_8 r^2 c + k_9 r c^2 + k_{10} c^3, \end{aligned}$$

where the center of the fitting neighborhood is taken as the origin. It quickly follows that the needed partials evaluated at local coordinates (r, c) are

$$\begin{aligned} \partial f/\partial r &= k_2 + 2k_4 r + k_5 c + 3k_7 r^2 + 2k_8 rc + k_9 c^2 \\ \partial f/\partial c &= k_3 + k_5 r + 2k_6 c + k_8 r^2 + 2k_9 rc + 3k_{10} c^2 \\ \partial^2 f/\partial r^2 &= 2k_4 + 6k_7 r + 2k_8 c \\ \partial^2 f/\partial c^2 &= 2k_6 + 2k_9 r + 6k_{10} c \\ \partial^2 f/\partial r\, \partial c &= k_5 + 2k_8 r + 2k_9 c. \end{aligned}$$

It is easy to see that if the above quantities are evaluated at the center of the pixel where local coordinates $(r, c) = (0, 0)$, only the constant terms will be of significance. If the partials need to be evaluated at an arbitrary point in a pixel's area, then a linear or quadratic polynomial value must be computed.

4.2. An Observation About Cubic Fits

A two-dimensional cubic polynomial includes an arbitrary quadratic polynomial, and thus features like pit, peak, and saddle can be replicated exactly. For other surface features like ridges or ravines, cubics are either exact or fairly decent approximations. It is

Fig. 4. Cubic fit of step edge that causes ravine and ridge to occur.

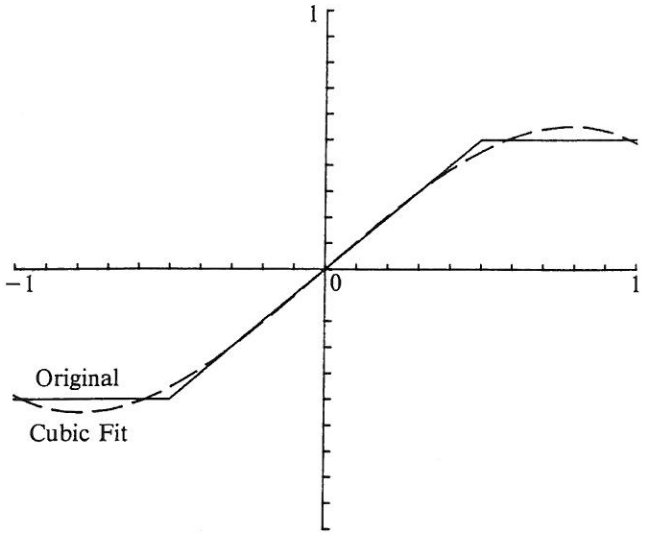

Fig. 5. Cubic fit (dashed line) of step edge (solid line) with inflection point outside window.

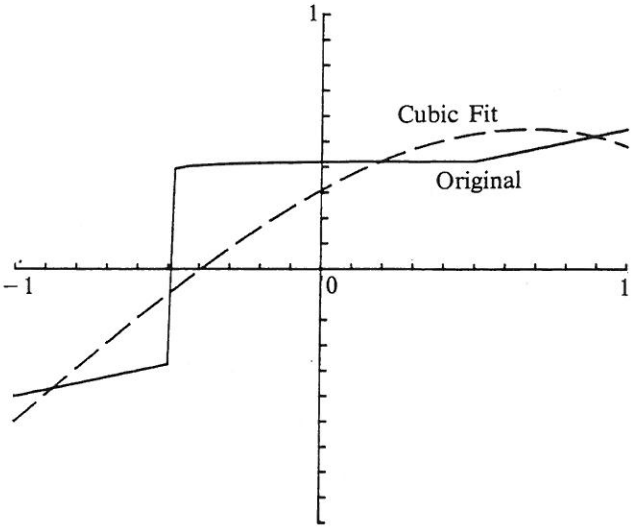

frequently possible to classify surface characteristics correctly even though the surface is not like any polynomial. For example, the center of the step edge shown in Fig. 4 can be accurately predicted by the inflection point of the cubic fit, which is quite good. However, there are smooth surfaces that (over the window being used) simply do not look like cubic polynomials, and feature classifications based on the best least-squares cubic polynomial approximation will be incorrect. For example, in Fig. 4, the cubic fit shows a prominent ravine at the foot of the slope, and the foot pixel would be (incorrectly) labeled a ravine pixel. Figure 5 shows a steplike edge whose cubic fit has an inflection point outside the entire window!

5. Summary of the Topographic Classification Scheme

The scheme is a parallel process for topographic classification of every pixel, which can be done in one pass through the image. At each pixel of the image, the following four steps need to be performed.

1. Calculate the fitting coefficients, k_1 through k_{10}, of a two-dimensional cubic polynomial in an n-by-n neighborhood around the pixel. These coefficients are easily computed by convolving the appropriate masks over the image.
2. Use the coefficients calculated in step 1 to find the gradient, gradient magnitude, and the eigenvalues and eigenvectors of the Hessian at the center of the pixel's neighborhood, (0, 0).
3. Search in the direction of the eigenvectors calculated in step 2 for a zero-crossing of the first directional derivative within the pixel's area. (If the eigenvalues of the Hessian are equal and nonzero, then search in the Newton direction.)
4. Recompute the gradient, gradient magnitude, and values of second directional derivative extrema at each zero-crossing. Then apply the labeling scheme as described in Sections 3.1–3.4.

6. Examples

In this section, we show the results of the topographic primal sketch on several test images, three of which are simply described mathematical surfaces

Fig. 6. A. Saddle surface.
B. Topographic labeling of saddle surface.

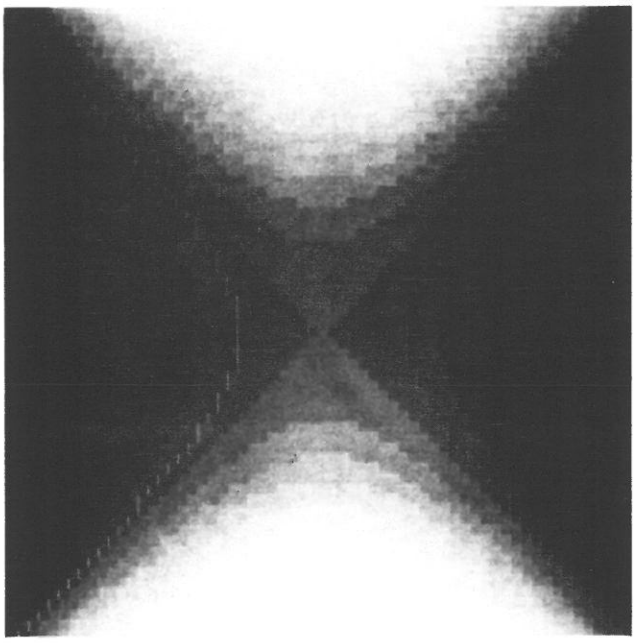

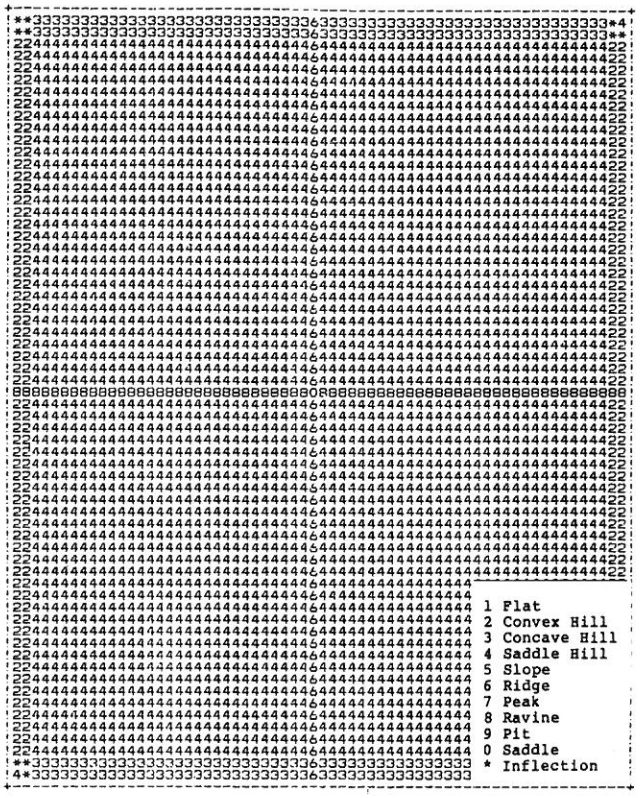

and two of which are real images. We will also examine how the size of the window affects the results of the classifier.

6.1. Saddle Surface

A perfect saddle surface of size 64 × 64 having no noise can be generated by the equation

$$f(r, c) = r^2 - c^2.$$

Taking the origin at image coordinates (32, 33), the surface plot is as illustrated in Fig. 2, and the gray-level image of the saddle surface is as illustrated in Fig. 6A. The results of the classifier are shown in Fig. 6B. Each number in the figure represents the label assigned the pixel by the classifier. As expected, a ridge-line one pixel in width was found running north-south, and, orthogonal to the ridge-line, a ravine-line one pixel in width was found. The center pixel of the surface was correctly classified as a saddle point. All other pixels on the surface were correctly classified as saddle hillsides.

Next, we add Gaussian noise to the saddle surface. The noise has a mean of 0.0 and standard deviation of 4.0. The results of the classifier using different sized windows (5 × 5, 7 × 7, 9 × 9, 11 × 11) on the noisy surface are shown in Fig. 7. As the window size increases, the results of the classifier improve dramatically. The classification resulting from the 11 × 11 window is almost identical to the classification done on the original, perfectly smooth surface. This would seem to suggest using as large a window size as possible, but in the next example we will show that this is not always a good idea.

6.2. Ridges and Valleys

A series of ridges and valleys can be generated across the column direction by the following equation:

Fig. 7. Neighborhood topographic labeling of noisy saddle surface showing ridge (black) and ravine (white). A. 5 × 5 window. B. 7 × 7 window. C. 9 × 9 window. D. 11 × 11 window.

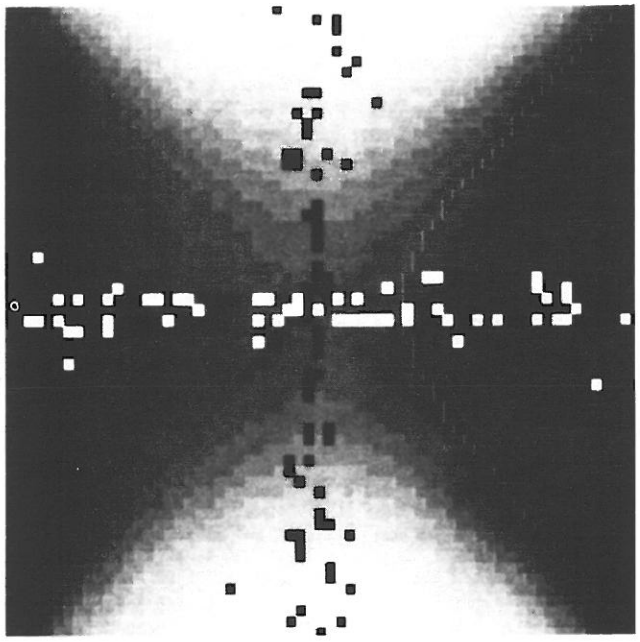

7a

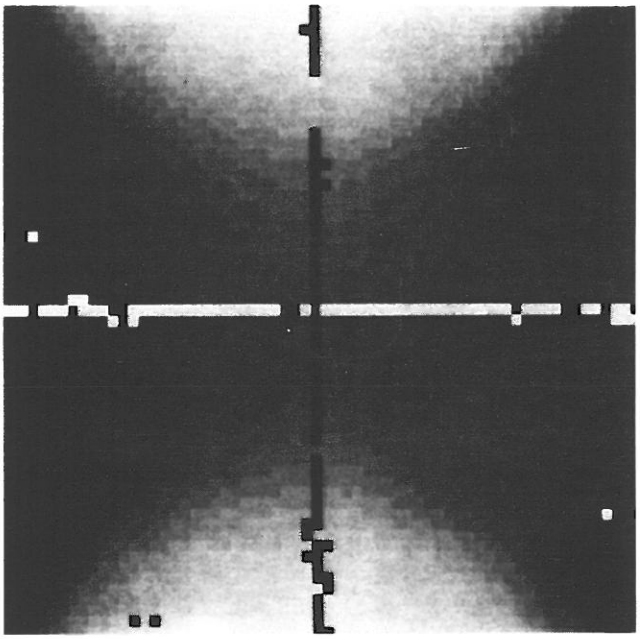

7c

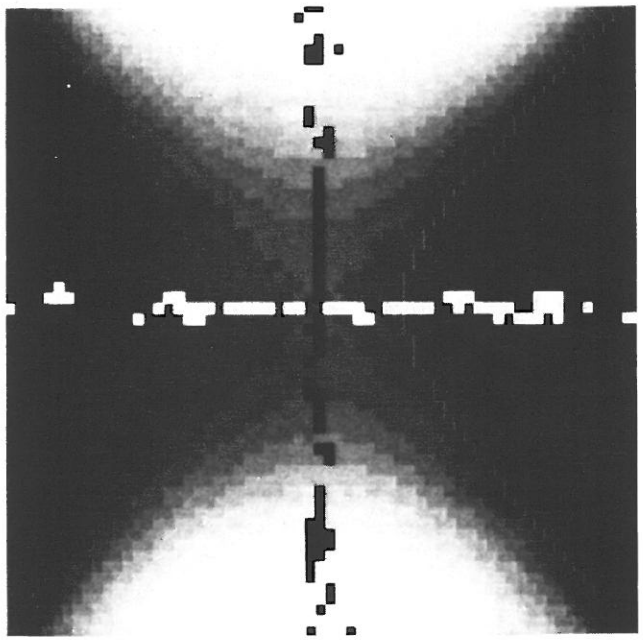

7b

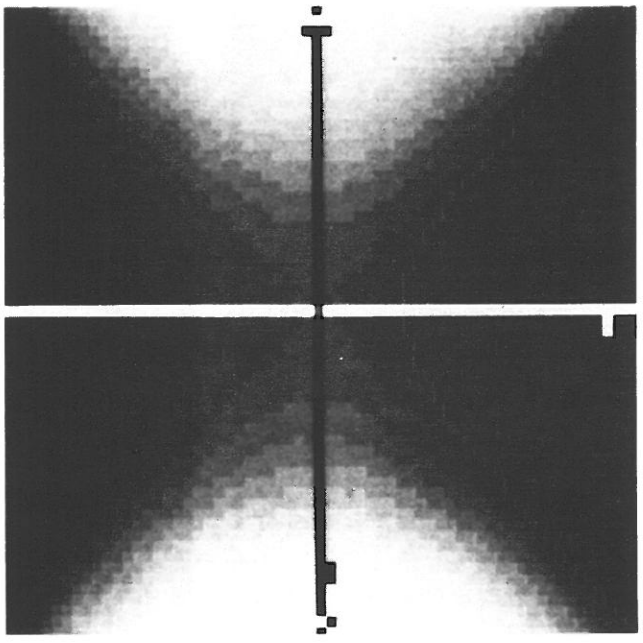

7d

Fig. 8. Image and topographic labeling of sine waves. A. Image. B. Topographic labeling with 5 × 5 window. C. Labeling with 9 × 9 window. D. Labeling with 13 × 13 window.

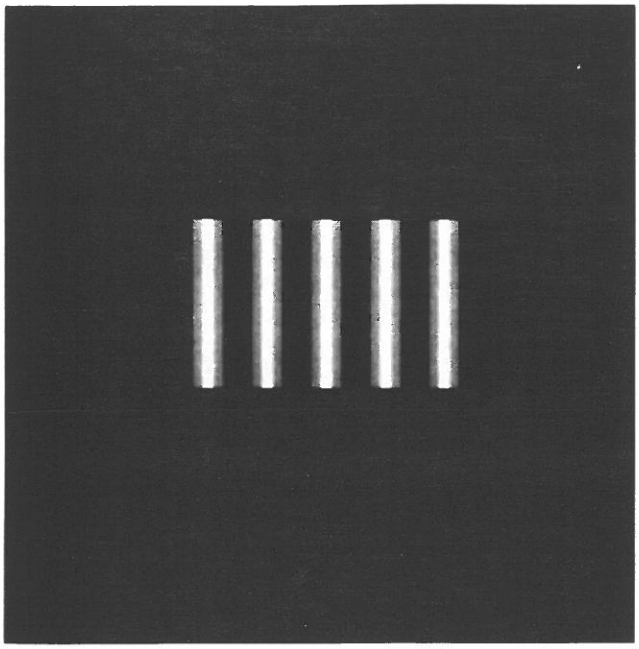

8a

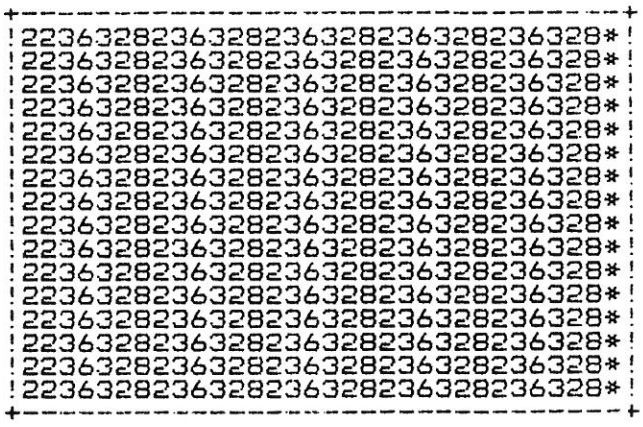

8b

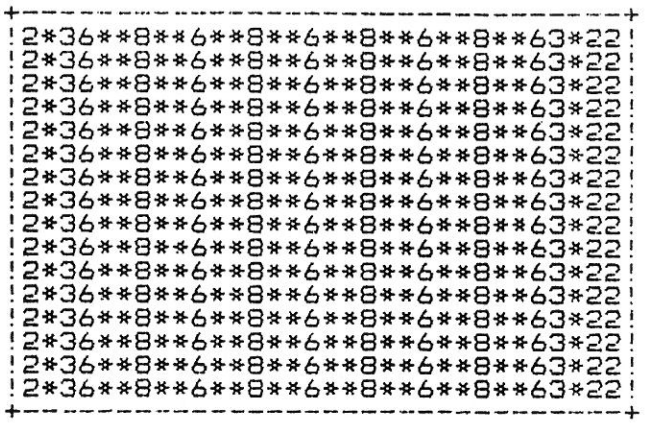

8c

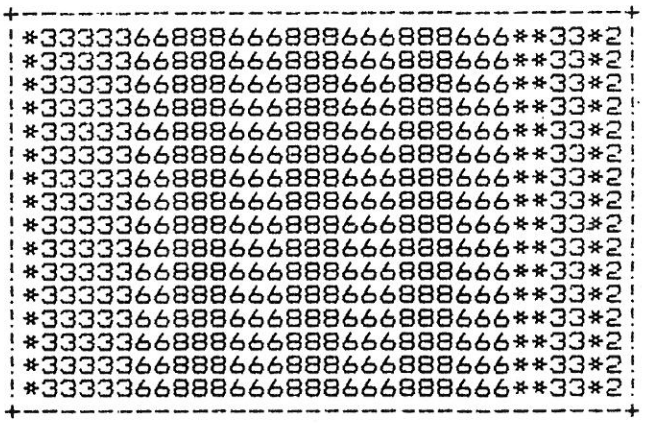

8d

$$f(r, c) = \begin{cases} 1 \text{ if col} = 1, 7 \\ 2 \text{ if col} = 2, 6 \\ 3 \text{ if col} = 3, 5 \\ 4 \text{ if col} = 4 \end{cases}$$

where col = mod(c − 1, 7) + 1.

It is easy to see from the above equation that every row will be the same with a ridge occurring every sixth pixel beginning in column 4, and a ravine occurring every sixth pixel beginning in column 7. The gray level plot of the 16 × 32 image is shown in Fig. 8A. As expected, the ridges and ravines are correctly identified. The results of the classifier using window sizes 5 × 5, 9 × 9, and 13 × 13 are found in Figs. 8B, C, and D respectively. As the window size increases, the results of the classifier become less accurate. This result is exactly the opposite of what happened on the saddle surface. The reason is that this surface is much more "busy" than the saddle surface. The larger window size on this particular surface results in too many complexities for the cubic fit to handle.

The conclusion is that the window size used should be a function of the noise and the complexity of the image surface. One should use as big a window size as possible without allowing the complexity of the

Fig. 9. Surface of revolution. A. Surface plot. B. Ravine ridge (black) and ravine (white) labeling of the surface.

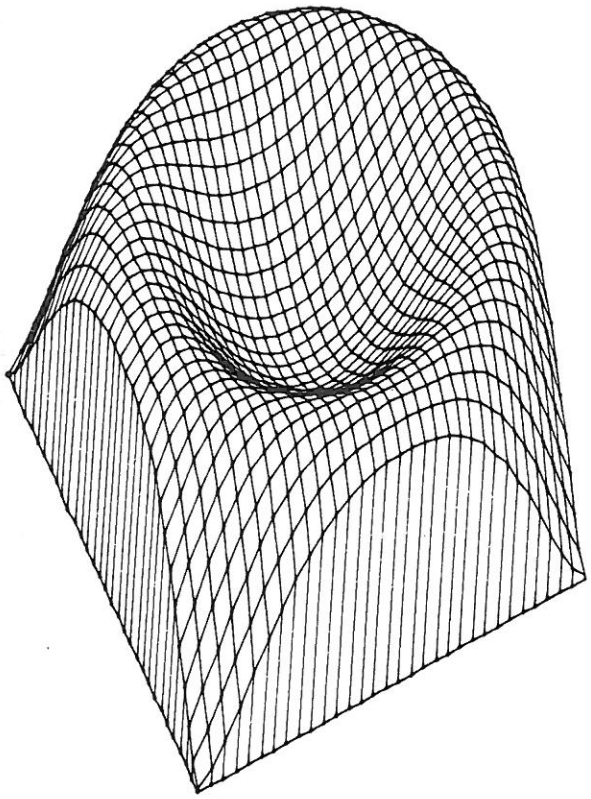

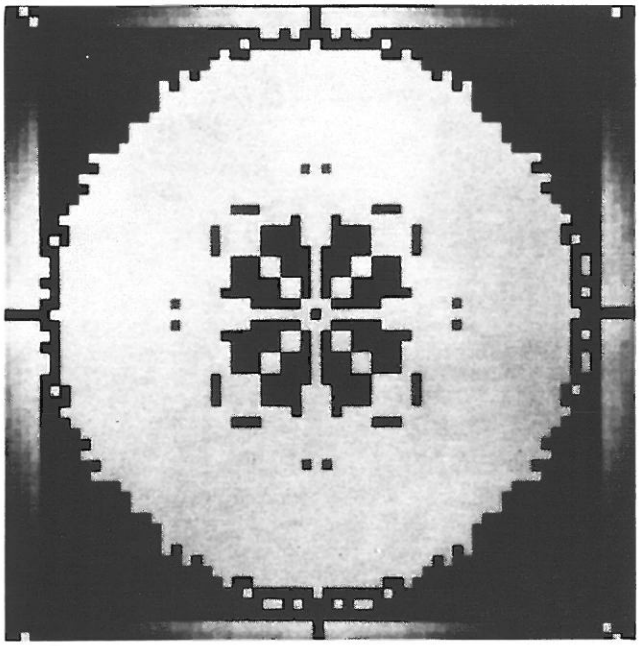

surface to degrade the cubic fit to any significant degree.

6.3. Surface of Revolution

A surface of revolution of size 64 × 64 with no noise can be generated by the equation

$$f(r, c) = k * \sin(0.5 * (r^2 + c^2)),$$

with origin at image coordinates (32, 32). The surface plot is illustrated in Fig. 9A. The topographic labeling on this surface shows some surprising results. A continuum of ridges and a continuum of ravines are found on the surface (see Fig. 9B). The reasons for these ridge and ravine points were discussed in Section 2.4. Also, the local cubic fits are very poor on this surface of revolution. This leads to some unexpected results, such as the peaks found on the rim of the surface where the ridges and ravines come together. The pixels labeled saddle on the image occur at locations where both a ridge and ravine were detected within the same pixel.

Notice that the labelings produced are not perfectly symmetric, as one would expect on a radially symmetric surface. The reason for this is that the cubic surface estimation is done with rectangular windows, which produces different cubic approximations at the same radial distance from the axis of revolution and hence radially unsymmetric labeling. Symmetric labeling would be produced by using a circular window, but choosing a particular window shape requires a priori knowledge of the nature of the image surface.

6.4. Real Image

In this section, we show the results of the classifier on two real images. The results on the top left corner of a chair image are illustrated in Fig. 10B. The results on the upper middle section of the bin of machine parts are illustrated in Fig. 11B. The vari-

Fig. 10. Results of the classifier on a real image. A. Chair. B. Upper left corner of chair. C. Ridges (black) and ravines (white). D. Hillside (white).

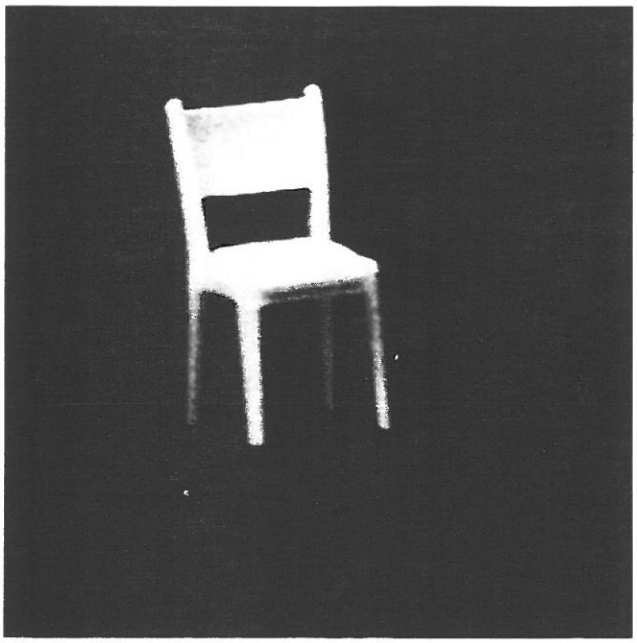

10a

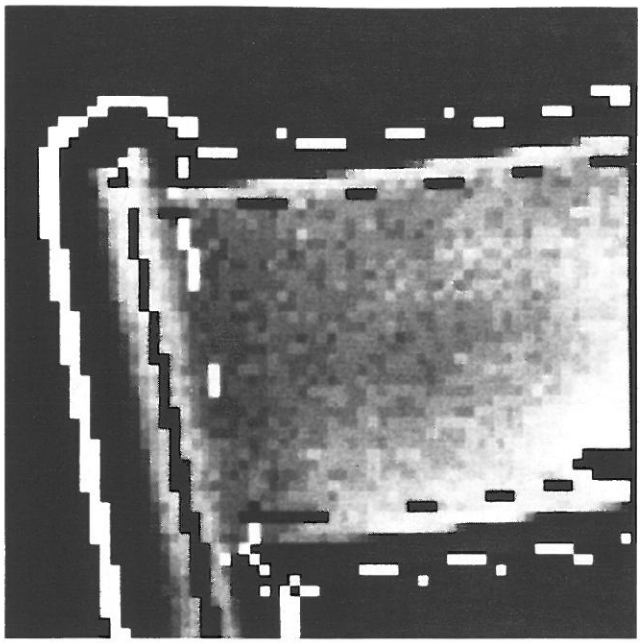

10c

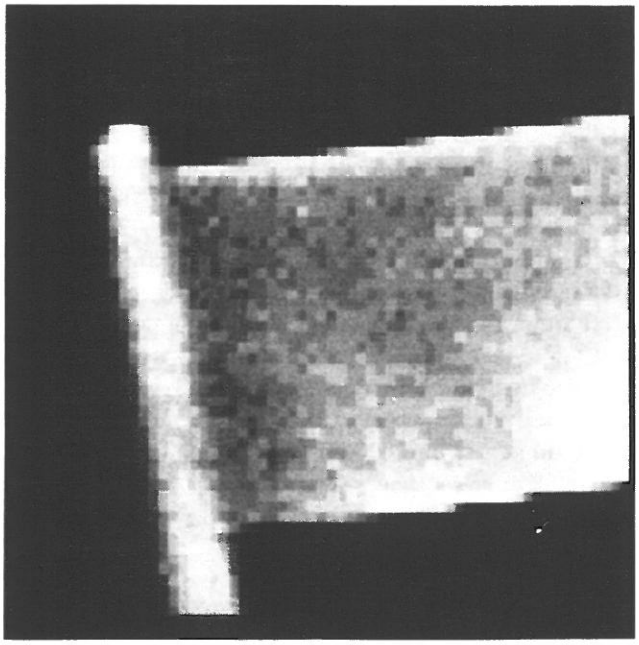

10b

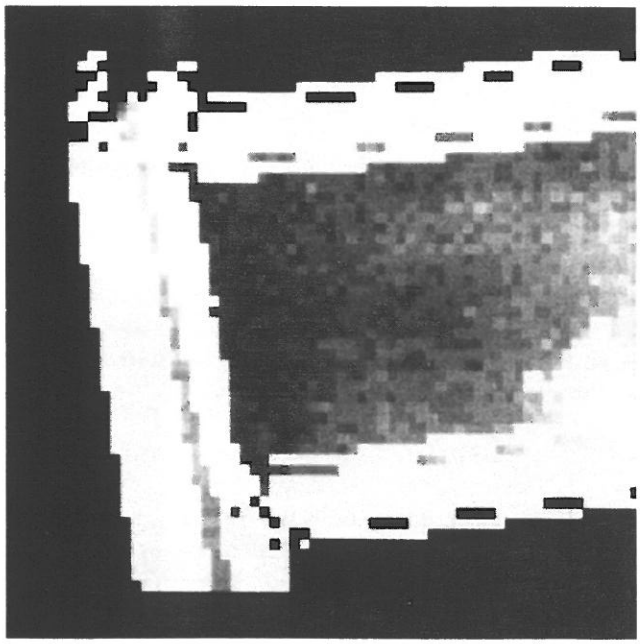

10d

Fig. 11. A. Screw. B. Ridges (black). C. Ravines (white). D. Convex hillside (white). E. Concave hillside (black).

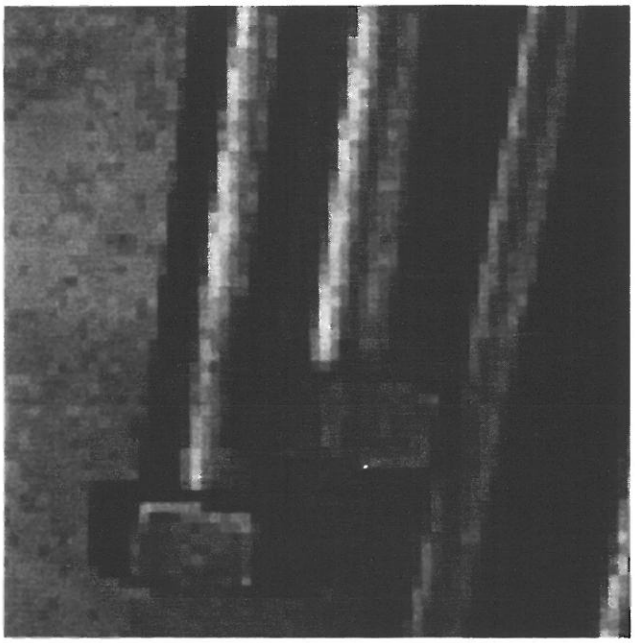

11a

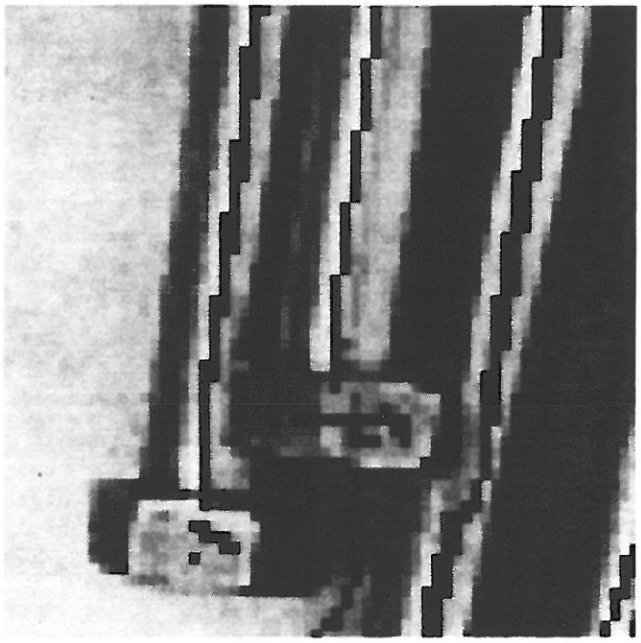

11b

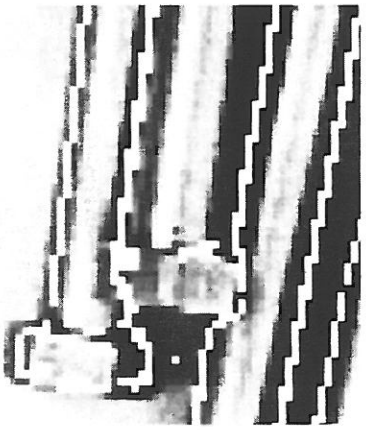

11c

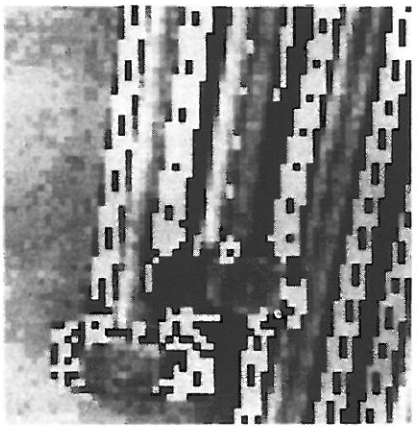

11d

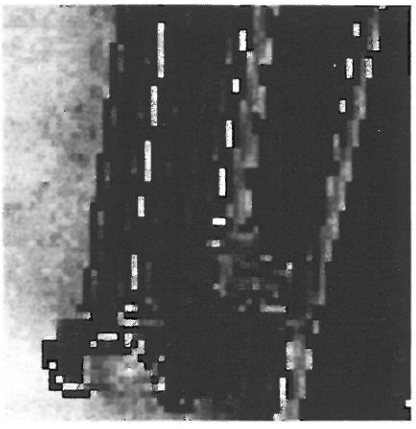

11e

Fig. 12. A. Machine parts. B. Upper left corner showing subimage ridges (black) and ravines (white). C. Center showing subimage ridges (black) and ravines (white).

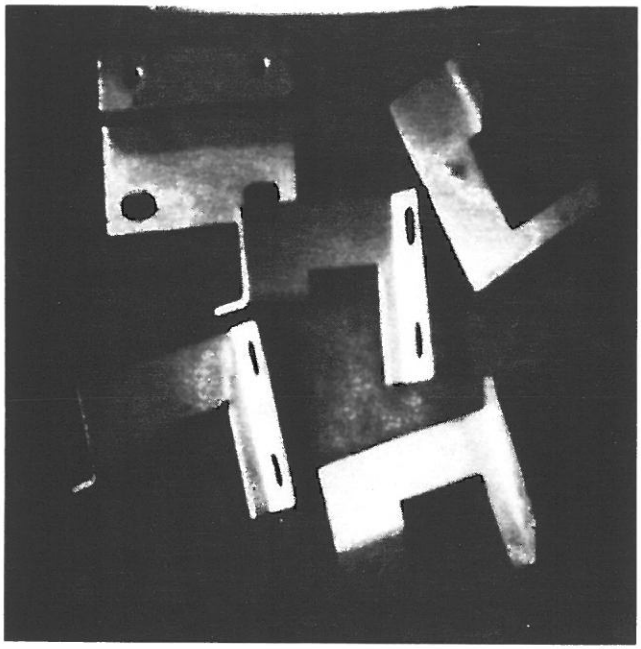

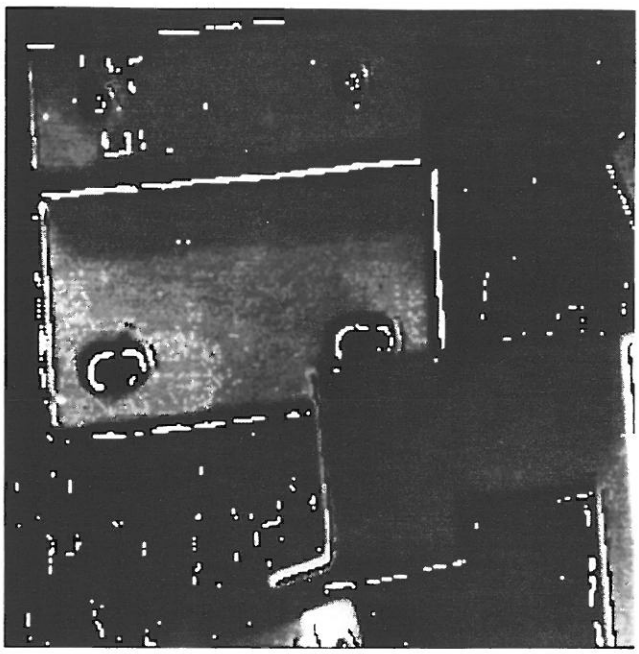

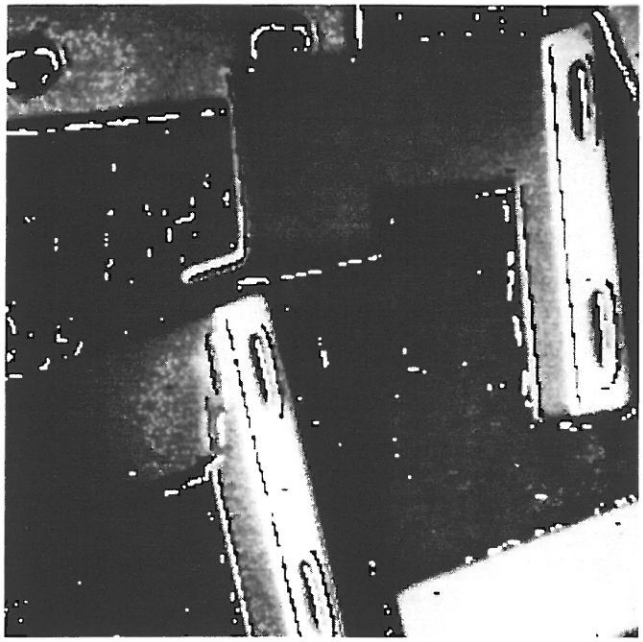

ous nonflat labels in the backgrounds of the images are caused by very slight dips and rises in the cubic surface fit. These may be cleaned up by requiring the eigenvalues to be above a certain threshold to be considered nonzero (see Fig. 10C). Figure 12 shows the labeling on an image of manufacturing parts. Notice how the highlighting can occur depending on the positioning of the parts. The ridge labels are quite useful for determining where the highlighting occurs.

7. Conclusions

In this paper, we have given a precise mathematical description of the various topographic structures that occur in a digital image and have called the classified image the topographic primal sketch. Our set of topographic categories is invariant under gray tone, monotonically increasing transformations and consists of peak, pit, ridge, ravine, saddle, flat, and hillside, with hillside being broken down further into the subcategories inflection point, slope, convex hill,

concave hill, and saddle hill. The hillside subcategories are not invariant under the monotonic transformations.

The topographic label assigned a pixel is based on the pixel's first- and second-order directional derivatives. We use a two-dimensional cubic polynomial fit based on the local facet model to estimate the directional derivatives of the underlying gray tone intensity surface. The calculation of the extrema of the second directional derivative can be done efficiently and stably by forming the Hessian matrix and calculating its eigenvalues and their associated eigenvectors. Strict, local, one-dimensional extrema (such as pit, peak, ridge, ravine, and saddle) are found by searching for a zero-crossing of the first directional derivative in the directions of extreme second directional derivative (the eigenvectors of the Hessian). We have also identified another direction of interest, the Newton direction, which points toward the extremum of a quadratic surface. The classification scheme was found to give satisfactory results on a number of test images.

7.1. Directions for Further Research

Further research on the topographic primal sketch needs to be done to (1) develop better basis functions, (2) make use of fitting error, (3) find a solution for the ridge (ravine) continuum problem, and (4) develop techniques for grouping of the topographic structures. Basis functions worth considering include trigonometric polynomials, polynomials of higher order, and piecewise polynomials of lower order than cubic. The basis functions problem is to find a set of basis functions and an associated inner product for least-squares approximation that can correctly replicate all common image surface features and be simultaneously computationally efficient and numerically stable.

Fitting error needs to be used in deciding into which class a pixel falls. Noise causes the fitting error to increase, and increased fitting error increases the uncertainty of the labeling. Also, global knowledge of how the topographic structures fit together could be used to correct the misclassification error caused by noise. The way the neighborhood size affects the surface fitting error and the classification scheme needs to be investigated in detail.

The ridge (ravine) continuum problem needs to be solved. It may be that there is no way to distinguish between a true ridge and a ridge continuum using only the values of partial derivatives at a point. The solution may require complete use of the partial derivatives in a local area about the pixel.

Most important for the use of the primal sketch in a general robotics computer vision system is the development of techniques for grouping and assembling topographically labeled pixels to form the primitive structures involved in higher-level matching and correspondence processes. How well can stereo correspondence or frame-to-frame time-varying image correspondence tasks be accomplished using the primitive structures in the topographic primal sketch? How effectively can the topographic sketch be used in undoing the confounding effects of shading and shadowing? How well will the primitive structures in the topographic sketch perform in the two-dimensional to three-dimensional object-matching process?

References

Ehrich, R. W., and Foith, J. P. 1978. Topology and semantics of intensity arrays. *Computer vision systems*, New York: Academic, pp. 111–128.

Grender, G. C. 1976. TOPO III: A Fortran program for terrain analysis. *Comput. Geosci.* 2:195–209.

Haralick, R. M. 1980. Edge and region analysis for digital image data. *Comput. Graphics Image Processing* 12(1): 60–73.

Haralick, R. M. 1981. The digital edge. *Proc. 1981 Conf. Pattern Recognition Image Processing*. New York: IEEE Computer Society, pp. 285–294.

Haralick, R. M. 1982. Zero-crossing of second directional derivative edge operator. *SPIE Proc. Robot Vision*. Bellingham, Wa.: SPIE.

Hsu, S., Mundy, J. L., and Beaudet, P. R. 1978. Web representation of image data. *4th Int. Joint Conf. Pattern Recognition*. New York: IEEE Computer Society, pp. 675–680.

Johnston, E. G., and Rosenfeld, A. 1975. Digital detection of pits, peaks, ridges, and ravines. *IEEE Trans. Syst. Man Cybern.*, July, pp. 472–480.

Laffey, T. J., Haralick, R. M., and Watson, L. T. 1982. Topographic classification of digital image intensity surfaces. *Proc. IEEE Workshop Comput. Vision: Theory Contr.* New York: IEEE Computer Society, pp. 171–177.

Lee, H. C., and Fu, K. S. 1981. The GLGS image representation and its application to preliminary segmentation and pre-attentive visual search. *Proc. 1981 Conf. Pattern Recognition Image Processing.* New York: IEEE Computer Society, pp. 256–261.

Marr, D. 1976. Early processing of visual information. *Philosophical Trans. Royal Soc. London* B 275:483–524.

Marr, D. 1980. Visual information processing: The structure and creation of visual representations. *Philosophical Trans. Royal Soc. London* B 290:199–218.

Paton, K. 1975. Picture description using Legendre polynomials. *Comput. Graphics Image Processing* 4(1): 40–54.

Peuker, T. K., and Johnston, E. G. 1972 (Nov.). Detection of surface-specific points by local parallel processing of discrete terrain elevation data. Tech. Rept. 206. College Park Md.: University of Maryland Computer Science Center.

Peuker, T. K., and Douglas, D. H. 1975. Detection of surface-specific points by local parallel processing of discrete terrain elevation data. *Comput. Graphics Image Processing* 4(4):375–387.

Rutishauser, H. 1971. Jacobi method for real symmetric matrix. *Handbook for automatic computation, volume II, linear algebra, ed.* J. H. Wilkinson and C. Reinsch. New York: Springer-Verlag.

Strang, G. 1980. *Linear algebra and its applications,* 2nd ed. New York: Academic, pp. 243–249.

Toriwaki, J., and Fukumura, T. 1978. Extraction of structural information from grey pictures. *Comput. Graphics Image Processing* 7(1):30–51.