

A STATISTICAL TECHNIQUE FOR RECOVERING SURFACE ORIENTATION FROM TEXTURE IN NATURAL IMAGERY

Andrew P. Witkin

Artificial Intelligence Center
SRI International, Menlo Park, CA 94025

ABSTRACT

A statistical method is reported for inferring the shape and orientation of irregularly marked surfaces using image geometry. The basis for solving this problem lies in an understanding of projective geometry, coupled with simple statistical models of the contour generating process. This approach is first applied to the special case of surfaces known to be planar. The distortion of contour shape imposed by projection is treated as a signal to be estimated, and variations of non-projective origin are treated as noise. The resulting method is next extended to the estimation of curved surfaces, and applied successfully to natural images. The statistical estimation strategy is then experimentally compared to human perception of orientation: human observers' judgements of tilt correspond closely to the estimates produced by the planar strategy.

I INTRODUCTION

Projective geometry lawfully relates the shape of a surface marking, the orientation of the surface on which the marking lies, and the shape of the marking's projection in the image: given any two, the third can usually be recovered. This paper addresses the problem of recovering the shape and orientation of an irregularly marked surface from the shapes of the projected markings. Of the three constituents of the projective relation, one--projected shape--is given, and another--surface orientation--is to be recovered. Since two of the constituents must be known to recover the third, the problem has no solution unless something is assumed about the unprojected shapes of the surface markings. For instance, if those shapes were known exactly, recovering surface orientation would usually be a straightforward exercise in projective geometry.

More interesting and more general is the case of a surface about which nothing specific is known in advance. To recover surface orientation in this case, some assumption must be made about the geometry of surface markings that is general enough to apply to a broad range of surfaces, powerful enough to determine a solution for surface orientation, and true enough to determine the right solution.

To meet these requirements, the inference of surface shape will be treated as a problem of statistical estimation, combining constraints from projective geometry with simple statistical models of the processes by which surface markings are formed. The distortion imposed on shapes by projection will be treated, quite literally, as a signal, and the shapes themselves, as noise. Both the "signal" and the "noise" contribute to the geometry of the image, and statistical models of the noise permit the projective component to be isolated.

II ESTIMATING THE ORIENTATION OF PLANAR SURFACES

The estimation problem will first be considered subject to the artificial restriction that the surface is known to be planar. While not realistic, this limited case provides the groundwork from which more general methods will be developed.

The shape of a surface marking is the shape of its bounding curve. The shapes of curves are naturally described by tangent direction as a function of position. Since tangent direction in the image is subject to projective distortion, this representation is suitable for estimating projective distortion. The mapping between a tangent direction on a surface and the corresponding direction in the image is readily expressed as a function of surface orientation, by means of simple projective geometry.

By examining a region of the image, a distribution of projected tangent directions may be collected. The form of this distribution depends in part on the distribution of directions over the corresponding surface region, but undergoes systematic distortion depending on the orientation of the surface: the projection of an inclined shape is foreshortened, i.e. compressed in the direction of steepest inclination (the tilt direction.) The amount of compression varies with the angle between the image plane and the plane of the shape (the slant angle.)

The effect of this distortion on the distribution of tangent directions may be illustrated by the simple case of a circular marking. The distribution of tangent directions on the original circle, measured by arc length, is

uniform. The orthographic projection of a circle is an ellipse, whose minor axis lies parallel to the tilt direction, and whose eccentricity varies with the slant angle. The distribution of tangent directions on an ellipse is not uniform, but assumes minima and maxima in the directions of the minor and major axes respectively. The degree of nonuniformity increases with the eccentricity of the ellipse.

In other words, projection systematically "pushes" the image tangents away from the direction of tilt. The greater the slant angle, the more the tangents are "pushed." This systematic distortion of the tangent distribution is the projective "signal" that encodes surface orientation. Its "phase" (the direction toward which the tangents gravitate) varies with surface tilt, and its "amplitude" (the amount of distortion,) with surface slant. To estimate the phase and amplitude of the projective signal is to estimate surface orientation.

Statistical estimation of the projective component requires a model of the "noise", i.e. of the expected distribution of tangent directions prior to projective distortion. With no prior knowledge of the surface, there is no reason to expect any one tangent direction to be more likely than any other. In other words, it is natural to assume that all tangent directions on the surface are equally likely. (Note that the amount of projective distortion increases with surface slant, effectively increasing the signal-to-noise ratio. Therefore, as slant increases, the exact form assumed for the noise distribution becomes less critical.)

Together with the geometric relation, this simple statistical model defines a probability density function for surface orientation, given a set of tangent directions measured in the image. The surface orientation value at which this function assumes a maximum is the maximum likelihood estimate for surface orientation, given the model. The integral of the function over a range of surface orientations is the probability that the actual orientation lies in that range. Conceptually, the distribution of tangent directions in the image can be projected onto an arbitrarily oriented planar surface. We seek that orientation for which the distribution of projected tangent directions is most nearly isotropic.

The estimator was first applied to geographic contours: projections of coastlines drawn from a digitized world map. This choice of data circumvents the problem of contour detection, and allows the actual orientation to be precisely controlled. The overall accord between estimated and actual orientation was excellent, and, equally important, the confidence measures generated by the estimator effectively distinguished the accurate estimates from the inaccurate ones.

The same technique was then applied to natural images, using zero-crossing contours in the convolution of the image with the Laplacian of a Gaussian [1], [2]. While the veridical orientations were not independently measured, the maximum likelihood estimates are in close accord with the perceived orientations (see Figure 1).

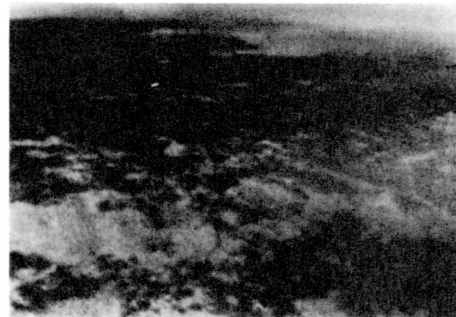
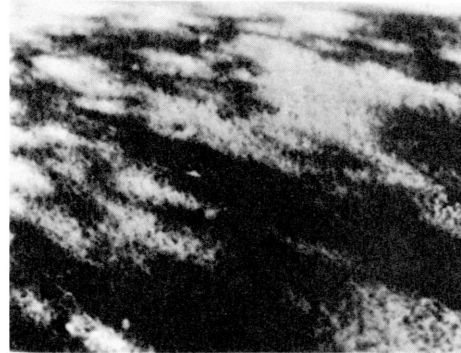


Figure 1 - Two photographs of roughly planar surfaces, and the orientation estimates obtained from them. The estimated orientations are indicated by ellipses, representing the projected appearance a circle lying on the surface would have if the maximum likelihood estimate were correct.

III EXTENTION TO CURVED SURFACES

To apply the methods developed in the planar case to curved surfaces without additional assumptions, it would be necessary to obtain at each point in the image a measure of the distribution of tangent directions. But such a local measure is never available, because the density of the contour data is limited. On the other hand, a distribution can be taken at each point of the data in a surrounding region, as small as possible, but large enough to provide a reasonable sample. This spatially extended distribution may be represented as a three dimensional convolution of the image data with a summation function.

To understand how such a distribution should be applied to estimate surface disposition, it is helpful to distinguish the intuitive, perceptual notion of surface orientation from the strict definition of differential geometry. It may be argued that surface orientation is not a unique

property of the surface, but must be regarded as a function of scale. The scale at which orientation is described corresponds to the spatial extent over which it is measured. Thus, by measuring orientation over a large extent, the surface is described at a coarse scale. The scale at which the surface is estimated therefore depends on the spatial extent over which the distribution is computed. Since that extent must be sufficiently large compared to the density of the data, the density effectively limits the spatial resolution of the estimate. On this view, close parallels can be drawn to Horn's [3] method for inferring shape from shading. The local estimator for orientation is a geometric analogue to the photometric reflectivity function.

The strategy was implemented, and applied to natural images. Contours were extracted as in the planar case, and the spatially extended distribution approximated by a series of two-dimensional convolutions with a "pillbox" mask. The estimated surfaces were in close accord with those perceived by the human observer (see Figure 2).

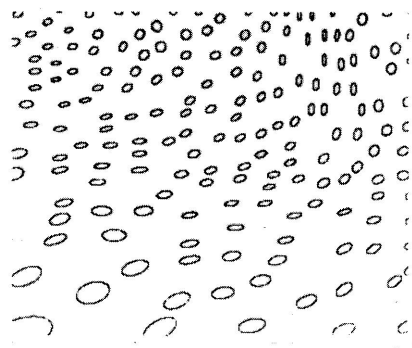
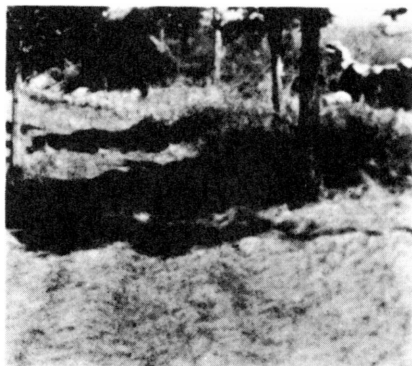


Figure 2 - A complex image, and the estimate obtained. The ellipses represent the appearance the estimated surface would have were it covered with circles. Note that the estimate correctly distinguishes the highly slanted foreground from the nearly frontal background. The upward pitch of the right foreground is also detected.

IV RELATION TO HUMAN PERCEPTION

A psychophysical experiment was performed to examine the relation between human observers' judgments of the orientations of curves perceived as planar, and the estimates obtained from the estimation strategy outlined above.

A series of "random" curves were generated using a function with pseudorandom parameters. Although such curves have no "real" orientation outside the picture plane, they often appear inclined in space. Observer's judgments of orientation were obtained by matching to a simple probe shape. The judgments of tilt (direction of steepest descent from the viewer) were highly consistent across observers, while the slant judgments (rate of descent) were much more variable.

Orientation estimates for the same shapes were computed using the planar estimator, and these estimates proved to be in close accord with those of the human observers, although the shapes had no "real" orientation.

While no conclusion may be drawn about the mechanism by which human observers judge orientation from contours, or about the measures they take on the image, this result provides evidence that the human strategy, and the one developed on geometric and statistical grounds, are at the least close computational relatives.

REFERENCES

1. Marr, D. C & Poggio, T., "A computational theory of human stereo vision," Proc. Roy. Soc. Lond., Vol. 204, pp.301-328 (1979)
2. Marr, D. C & Hildreth, E., "A Theory of edge detection," MIT AI Memo 518 (1979)
3. Horn, B. K. P., "Understanding image intensities," Artificial Intelligence, Vol. 21, No. 11, pp.201-231 (1977)